

Random Quotients of Hyperbolic Spaces with Respect to WPD Actions and Linear Connectivity of the Boundary

Inhyeok Choi and Henry Talbott

October 15, 2026

Abstract

We study random quotients of acylindrically hyperbolic groups and hierarchically hyperbolic groups, whose (partial) hyperbolicity is captured by their actions on Gromov hyperbolic spaces. We show that if an acylindrically hyperbolic group G acts on a Gromov hyperbolic space X coarsely transitively with a WPD element, then the quotient of X by the normal closure of a random element of G has linearly connected boundary. If G is a group with a (relative) hierarchically hyperbolic group structure, with a quasi-tree top domain X , then in fact the quotient of X by the normal closure of a random element of G has boundary homeomorphic to the Nöbeling curve.

Our method relies on a careful study of the configuration of translates of typical periodic geodesics in X .

1 Introduction

We consider *few-relators random quotients* of groups acting on hyperbolic spaces by isometries. We are especially interested in the action $G \curvearrowright X$ of a *hierarchically hyperbolic group* (HHG) on the maximal domain of the corresponding hierarchically hyperbolic space (HHS) structure, such as a mapping class group acting on its curve graph, or a right-angled Artin group (RAAG) acting on its coned-off contact graph. Our model of random quotients of HHGs is largely equivalent to the model defined by Abbott-Berlyne-Mangioni-Ng-Rasmussen, for which the authors obtained the result that random quotients of HHGs are hierarchically hyperbolic [Abb+26].

Specifically, let $G \curvearrowright X$ be a cobounded, non-elementary, WPD, and *action path-connected* (see Definition 2.27) group action by isometries on a Gromov hyperbolic space; WPD and action-path-connectedness are immediate in the HHG setting. Given admissible random walks $(w_n^{(1)}, \dots, w_n^{(k)})_{n \geq 0}$ on G , we draw the n -step elements, define

$$N := \langle\langle w_n^{(1)}, \dots, w_n^{(k)} \rangle\rangle,$$

and consider the induced action $G/N \curvearrowright X/N$. By the aforementioned work, X/N is Gromov hyperbolic, so we may consider the geometry of its boundary equipped with the visual metric.

We prove the following results. All asymptotic probabilities are taken with respect to the limit $n \rightarrow \infty$. Recall that a space is linearly connected if points may be joined by paths with diameters coarsely bounded above by the distances between those points.

Theorem 1.1. *Let $G/N \curvearrowright X/N$ be as above. Then a.a.s. $\partial(X/N)$ with its intrinsic visual metric is linearly connected.*

While many known examples of HHGs have quasi-tree maximal domains (e.g. cubical HHGs such as RAAGs) [BHS17], the canonical maximal domains for mapping class groups – the curve graphs – are un-tree-like in the sense of possessing highly connected boundaries [Gab09]. We find this dichotomy striking, especially in light of the role of mapping class groups as prototypical HHGs. For general acylindrically hyperbolic groups (a class including all HHGs), the foundational work of Bestvina-Bromberg-Fujiwara on projection complexes gives a general recipe for constructing cobounded WPD actions on quasi-trees for any such group [BBF15]. In contrast, to the best of the authors’ knowledge, we lack a general recipe to construct cobounded WPD actions on hyperbolic spaces that are not quasi-trees.

Theorem 1.1 shows that acylindrically hyperbolic groups acting on hyperbolic spaces that are not quasi-trees are highly abundant. In fact, such actions may be obtained by applying a quotient to *any* action of an HHG on its maximal domain, so long as the quotient is sufficiently ‘generic’ or ‘random’. As a consequence, we obtain a recipe for constructing many HHGs with maximal domains that are not quasi-trees.

In some cases, we can pin down the topology of the quotient boundary more precisely. We first obtain a general result controlling the dimension of the boundary.

Theorem 1.2. *Let $G/N \curvearrowright X/N$ be as above. Then a.a.s.*

$$1 \leq \dim(\partial(X/N)) \leq \dim(\partial X) + 1.$$

Note that the lower bound follows immediately from Theorem 1.1.

For our next result, note that the *dense unbounded domains* property required below (see Definition 9.6) is expected to hold for most relatively hierarchically hyperbolic groups (RHHGs). In particular, it holds for mapping class groups, for any right-angled Artin group with a defining graph that is connected and not a join, and many right-angled Coxeter groups. (Although note that Theorem 1.3 does not itself apply to mapping class groups due to the quasi-tree assumption.)

Theorem 1.3. *Let G/N and X/N be as above. Additionally, assume that $G \curvearrowright X$ is the action of a relatively hierarchically hyperbolic group on its maximal domain, that the RHHG structure for G satisfies the dense unbounded domains property, and that X is a quasi-tree. Then a.a.s. $\partial(X/N)$ is homeomorphic to the Nöbeling curve.*

Recall that the *Nöbeling curve* is a topological space of dimension 1 that may be constructed by deleting from $\mathbb{R}^3$ all points with at least two rational coordinates. To give a sense for this space, note that despite having topological dimension 1, it remains connected after deleting any countable subset.

We believe Theorem 1.3 is particularly interesting in light of analogous results on curve graph boundaries, discussed below.

1.1 Applications

We identify two classes of (relatively) hierarchically hyperbolic groups satisfying the hypotheses of Theorem 1.3.

1. Graphs of groups with edges representing commutation, as defined in [BR22], satisfying certain simple conditions. This class includes all right-angled Artin groups G such that the defining graph of G is connected and is not a join, and many right-angled Coxeter groups.
2. Fundamental groups of non-geometric 3-dimensional graph manifolds, using the hierarchical hyperbolicity structure developed in [Hag+25].

Two examples are illustrated in Figure 1. In both examples, we assign a group to each vertex, take a free product of these groups, and then declare the relations that generators of groups corresponding to adjacent vertices commute. For the left pentagon graph, we assign $\mathbb{Z}$ to each vertex; this is a right-angled Artin group. For the right graph, we arrange five pentagon graphs in a super-pentagon, and add edges between all pairs of vertices in adjacent pentagons. We then assign the group $\mathbb{Z}/2\mathbb{Z}$ to each vertex; this is a right-angled Coxeter group.

1.2 Background

Many groups in geometric group theory exhibit hyperbolic behavior through actions on metric spaces; *hierarchical hyperbolicity* is a way of axiomatizing a specific strong form of hyperbolicity observed across many important groups [BHS17; BHS19], while acylindrical hyperbolicity likewise axiomatizes a somewhat weaker form observed in groups such as $\text{Out}(F_n)$ [Osi16]. The prototypical example of hierarchical hyperbolicity is the Masur-Minsky machine for mapping class groups, in which mapping classes may be described via coordinates on Gromov hyperbolic spaces, the *curve graphs*, obtained from subsurfaces [MM99; MM00]. Curve graphs have proved to be extremely useful tools in their own right [MM99; Min06; Raf07; BM08; Beh+12; BCM12; BBF15]. The HHG machinery has since seen application to many other notable groups, including extra-large Artin groups [HMS22] and (very recently) the fundamental group of the moduli space of smooth complex cubic surfaces [AW26].




Figure 1: Two defining graphs that provide examples of groups satisfying the hypotheses of Theorem 1.3. The left graph is taken as the defining graph of a right-angled Artin group. The right graph is constructed by taking five pentagon graphs, placing these graphs in a pentagonal configuration, and adding edges between all pairs of vertices in adjacent graphs. This graph is taken as the defining graph of a right-angled Coxeter group.

The Gromov boundary and its generalizations are central tools in the study of Gromov hyperbolic spaces, and are key to applications such as rigidity [CT07; BM91; Bow98; KK11; Kle07; CH17; DMS25]. The boundary of a curve graph is the corresponding space of *ending laminations* [Kla22]; boundary path connectivity of curve graphs was proven for special cases by Leininger-Schleimer and Leininger-Mj-Schleimer [LS09; LMS11], and in full generality by Gabai [Gab09]. Gabai further proved that curve graphs of punctured spheres have Nöbeling space boundaries [Gab09], generalizing prior work of Hensel-Przytycki [HP11]. In turn, Wright showed that all curve graphs have linearly connected boundaries [Wri24].

Gromov proposed random groups (specifically, random quotients of free groups) as models for the expected properties of ‘typical’ hyperbolic groups [Gro93]. Our results may be viewed as extensions of previous results on random hyperbolic groups to the HHG setting; Champetier showed few-relator random quotients of free groups have boundaries homeomorphic to the Menger curve [Cha95], while Dahmani-Guirardel-Przytycki showed the same for random quotients of free groups in the density model [DGP11]. We also note recent work of Ioannou on hyperbolic groups with Menger curve boundaries [Ioa26]. One may likewise study the statistical properties of random elements of groups; in our setting of acylindrically hyperbolic group actions $G \curvearrowright X$, we rely substantially on deep results on random walks due to Maher–Tiozzo [MT18; MT21].

1.3 Methods

Our methods have their earliest antecedents in the work of Champetier [Cha95], who studied random quotients of free groups and showed that these quotients have boundaries homeomorphic to the Menger curve (also called the Menger sponge). In order to adapt this approach to general acylindrically hyperbolic groups and hierarchically hyperbolic groups, one has to understand the structure of the normal closures of WPD elements. This structure was first systematically understood by Dahmani-Guirardel-Osin [DGO16], and further developed by Clay-Mangahas-Margalit [CMM21] and Clay-Mangahas [CM22]. These works introduce and develop the *rotating/spinning families* machinery, and while this machinery does not appear directly in our work, the core ideas underlie many of our arguments. Another important input is the projection complex machinery due to Bestvina-Bromberg-Fujiwara [BBF15], which was further refined by Bestvina-Bromberg-Fujiwara-Sisto [Bes+19].

To illustrate our methods for Theorem 1.1, we sketch the case of the free group. Let $2 \leq k \leq n$, let $G = \langle a_1, a_2, \dots, a_n \rangle$ be a free group of rank n , and let X be its standard Cayley graph.

Let us now consider a probability measure μ , the support of which is $\{a_1^\pm, \dots, a_n^\pm\}$. Then the random walk generated by μ is admissible with bounded steps on X . By the work of Maher-Tiozzo, there exists $\lambda > 0$ such that the following property holds for asymptotic probability 1: the element w obtained at step n of the random walk satisfies

$$0.9999\lambda n < d_X(e, w) < 1.0001\lambda n, \quad (w^{-1} \cdot w)_e \leq \log n.$$

Moreover, $[e, w]$ and $g[e, w]$ will overlap for length $< 0.0001n$ for each $g \in G \setminus \{e\}$.

Now consider the quotient $X/\langle\langle w \rangle\rangle$. What is the geometry of this space? It is no longer a tree, as each left coset of $\langle w \rangle$ gives rise to loops of length $\approx n$. However, $X/\langle\langle w \rangle\rangle$ retains Gromov hyperbolicity almost surely [Gro93]; this result in the general non-proper setting is studied in [Abb+26].

We now describe our proof of linear connectivity. The basic idea, as in [Cha95] and also [Wri24], is that given a path lying in the n -sphere S_n centered at some basepoint x_0 , we may ‘push it out’ to a finite-length path lying in $S_{n+1} \cup S_{n+2} \cup \dots \cup S_{n+K}$. Furthermore, we may do this pushing in a well-controlled way. Specifically, given a path $\gamma_n = (v_0, \dots, v_k)$ on the n -sphere S_n , for each $0 \leq i \leq k$ there exists a ‘prong outward’ at v_i . Namely, there exists a vertex w_i at distance from v_i such that the directed paths $\overrightarrow{v_i w_i}$ ’s are all directed away from the center x_0 and are repelling each other; in numeric terms, both $(w_i \cdot x_0)_{v_i}$ and $(w_i \cdot w_{i+1})_{v_i}$ are at most, say, 4. As w is a generic word, it contains all $w_i^{-1}(v_i^{-1}v_{i+1})w_{i+1}$ (of length ≤ 21) as subwords. Hence, in $X/\langle\langle w \rangle\rangle$, we have loops attached at each of $\{v_0, \dots, v_n\}$, where adjacent loops overlap for length 10. We assign the label u_i to the point at the apex of the i th loop, and join v_0 and v_k through a path γ_{n+1} that ‘skips’ along the loops, passing through each u_i . This path is illustrated in Figure 2. As the prongs are directed away from x_0 , this path lies in $S_{n+1} \cup S_{n+2} \cup \dots \cup S_{n+K}$. We note that the

overlapping construction of the loops is essential to keep the path from re-entering S_n ; as a particular consequence, the Gromov products between consecutive $w_1, \dots, w_m$ are strictly larger than the products between consecutive $v_1, \dots, v_k$.

We think of γ_{n+1} as a refined version of γ_n , and by inductively pushing paths γ_n deeper and deeper, we can draw limiting paths on $\partial X / \langle\langle w \rangle\rangle$ connecting arbitrary boundary points. While the pushing out procedure increases the length of the path at each step, we retain sufficient geometric control that the visual-metric diameter of the resultant boundary path is linearly bounded.

Our other results follow from further analysis of this basic scheme.




Figure 2: The heuristic idea for the inductive pushing out procedure. The pink curve represents the pushed-out path, which skips over the attached loops.

1.4 Acknowledgments

The authors would like to thank Alex Wright for suggesting the problem and Carolyn Abbott for further helpful discussions at the onset of the project. Furthermore, the authors would like to thank Giorgio Mangioni for his help with identifying applications of Theorem 1.3, in particular the 3-manifold group example. The authors would also like to thank George Domat and Alessandro Sisto for helpful conversations.

Talbott was supported by the National Science Foundation Graduate Research Fellowship under Grant. No. DGE 2241144. Choi was supported by the KIAS Individual

Grant (MG091901) at KIAS. Choi was also supported by the National Science Foundation under Grant No. DMS-2424139, while the authors were in residence at the Simons Laufer Mathematical Sciences Institute in Berkeley, California, during the Spring 2026 semester.

1.5 Artificial Intelligence Declaration

Talbott briefly experimented with using LLMs. The sole outcome of this usage was that the authors became aware of the Urysohn-Menger Sum Formula as a potential tool. Aside from this instance, no AI was used in the preparation of this paper. The text and images of the paper are entirely human-authored.

2 Preliminaries

2.1 Notation

Throughout the text, we will use the notation $A =_C B$ for real numbers A, B, C with $C \geq 0$ to denote $|A - B| \leq C$.

2.2 Hyperbolic Spaces and Projections

We review some hyperbolic geometry and introduce the key hypotheses we will need. All metric spaces we will consider in this paper are geodesic metric spaces, and in fact are simplicial graphs. As a consequence, the distance function is integer-valued. This assumption gives essentially no loss of generality, as whenever a group G acts on a metric space X by isometries, X is G -equivariantly quasi-isometric to a simplicial graph X_S , such that G also acts on X_S by isometries. We will further assume that G is finitely generated, so that G may be equipped with a word metric.

Given two points $x, y \in X$, we will use the notation $[x, y]_X$ to represent a geodesic from x to y . Recall that a geodesic metric space X is δ -hyperbolic (or *Gromov hyperbolic*) if for all geodesic triangles with sides $[x, y]_X$, $[y, z]_X$, and $[z, x]_X$, each side is contained in the δ -neighborhood of the union of the other two sides. It is often useful to consider *quasi-geodesics* instead of geodesics.

Throughout this section, we will assume a fixed δ -hyperbolic simplicial graph X .

Definition 2.1. A (λ, c) -quasi-geodesic is a map γ from a (possibly unbounded) subinterval I of $\mathbb{Z}$, such that for $s, t \in I$,

$$\frac{1}{\lambda} d_{\mathbb{Z}}(s, t) - c \leq d_X(\gamma(s), \gamma(t)) \leq \lambda d_{\mathbb{Z}}(s, t) + c$$

A c -coarse geodesic is a $(1, c)$ -quasi-geodesic.

A c -coarse geodesic at speed L is map γ from a subinterval $I \subseteq \mathbb{Z}$ to X such that for $s, t \in I$,

$$L d_{\mathbb{N}}(s, t) - c \leq d_X(\gamma(s), \gamma(t)) \leq L d_{\mathbb{N}}(s, t) + c$$

Remark 2.2. We will often use γ to refer interchangeably to a quasi-geodesic map and its image.

Definition 2.3. A subset $A \subseteq X$ is K -quasi-convex if for all $x, y \in A$ and all $[x, y]_X$, $[x, y]_X \subseteq \mathcal{N}_K(A)$.

We need a few standard lemmas and definitions.

Lemma 2.4 (Morse Lemma). *For all $\lambda \geq 0$ and $c \geq 0$, there is a constant $\Phi \geq 0$ dependent only on δ , λ , and c satisfying the following property: for any two (λ, c) -quasigeodesics γ_1 and γ_2 with the same endpoints in X , the Hausdorff distance between γ_1 and γ_2 is at most Φ .*

Definition 2.5 (Closest Point Projections). Assume $A \subseteq X$ is K -quasi-convex for some K . Given $x \in X$, we define

$$\pi_A(x) = \{y \in A : d(x, A) = d(x, y)\}.$$

For $B \subseteq X$, we define $\pi_A(B) = \bigcup_{x \in B} \pi_A(x)$. Lastly, for $A, B \subseteq X$, we define

$$d_A(B, C) = \text{diam}(\pi_A(B) \cup \pi_A(C))$$

Remark 2.6. Note that there are two quite different forms of projection that will be denoted with π : we will use π_A to denote closest-point projection in X to a quasi-convex subset of X , as well as $\pi_{X/N}$ to denote the quotient projection from X to X/N .

Lemma 2.7. *Assume $A \subseteq X$ is K -quasi-convex. Then there is a constant D dependent only on δ and K so that for any $x \in X$, $\text{diam}(\pi_A(x)) \leq D$.*

For the next lemma, the reader should note that the name ‘Bounded Geodesic Image’ is also frequently applied to a conceptually related but somewhat different axiom for hierarchically hyperbolic spaces.

Lemma 2.8 (Bounded Geodesic Image). *Let $A \subseteq X$ be K -quasi-convex, and fix $\lambda, c \geq 0$. Then there exists $B, B' > 0$ dependent only on δ , K , λ , and c so that if γ is a (λ, c) -quasi-geodesic disjoint from $\mathcal{N}_{B'}(A)$, then $\text{diam}(\pi_A(\gamma)) \leq B$.*

Lemma 2.9 (Large Projection Fellow-Traveling). *Let $A \subseteq X$ be K -quasi-convex. Then there exists Ω dependent only on δ and K so that for any points $x, y \in X$ and any $x' \in \pi_A(x)$, $y' \in \pi_A(y)$, if $d_A(x', y') \geq \Omega$, then any concatenation $[x, x']_X \cup [x', y']_X \cup [y', y]_X$ is a $(1, \Omega)$ -quasi-geodesic.*

In particular, $[x, x']_X \cup [x', y']_X \cup [y', y]_X$ and any $[x, y]_X$ must fellow-travel.

2.3 Gromov Products and the Boundary

Definition 2.10. Given $x, y, z \in X$, the *Gromov product* $(x \cdot y)_z$ is defined as

$$(x \cdot y)_z = \frac{d_X(x, z) + d_X(y, z) - d_X(x, y)}{2}.$$

Note that this definition makes sense in any metric space, although it is typically most useful in Gromov hyperbolic spaces.

The following observations are elementary.

Lemma 2.11. *Let x, y, z, w be points in X . Then the following properties hold:*

1. *We have*

$$d(z, [x, y]) - 2\delta \leq (x \cdot y)_z \leq d(z, [x, y]).$$

2. *We have*

$$(x \cdot y)_z + (z \cdot y)_x = d_X(x, y)$$

3. *(Triangle inequality for Gromov products.) We have*

$$(x \cdot y)_z \geq \min((x \cdot w)_z, (w \cdot y)_z) - 10\delta$$

Definition 2.12. Let $(x_0, x_1, \dots, x_n)$ be a sequence of points in X . We say these points are *C-aligned* for some constant $C > 0$ if for all $1 \leq i \leq n-1$,

$$(x_{i-1} \cdot x_{i+1})_{x_i} < C$$

Lemma 2.13 ([Gou22a, Lemma 3.7]). *Let $(x_0, \dots, x_n)$ be a sequence of C-aligned points in X such that*

$$d_X(x_i, x_{i+1}) \geq 2C + 2\delta + 1 \quad (i = 1, \dots, n-2).$$

Then any subsequence $(x_{i_1}, \dots, x_{i_m})$ of $(x_0, \dots, x_n)$ is $(C + \delta)$ -aligned. In particular, every geodesic $[x_0, x_n]_X$ passes through the $(C + 3\delta)$ -neighborhood of x_i for each $i = 1, \dots, n-1$.

We now turn to the Gromov boundary. Note that this definition does not in fact depend on the choice of basepoint.

Definition 2.14. Fixing a basepoint $x_0 \in X$, the *Gromov boundary* ∂X of X is the set of all sequences of points (x_i) such that $\liminf_{i,j} (x_i \cdot x_j)_{x_0} = \infty$, up to the equivalence relation that $(x_i) \sim (y_i)$ if $\lim_{i \rightarrow \infty} (x_i \cdot y_i)_{x_0} = \infty$.

The Gromov boundary carries a natural topology defined via *neighborhood bases* of points. Specifically, for all $\xi \in \partial X$ and $r \geq 0$, define $U(\xi, r)$ to be the set of all $\chi \in \partial X$ such that there exist sequences $(x_i), (y_i)$ with (x_i) in the equivalence class defining ξ , (y_i) in the equivalence class defining χ , and $\liminf_{i,j \rightarrow \infty} (x_i \cdot y_j)_{x_0} \geq r$. The collection of all $U(\xi, r)$ forms a neighborhood system at ξ .

Remark 2.15. Note that the sets $U(\xi, r)$ are not open in general under the neighborhood basis topology.

Lemma 2.16. *If X is proper, then for any $\xi \in \partial X$, there exists a ray $\gamma : \mathbb{N}_{\geq 0} \rightarrow X$ such that $(\gamma(i))$ lies in the equivalence class of ξ . For general X , there exists a $(1, 10\delta)$ -quasigeodesic ray γ with this property.*

The boundary ∂X is in fact metrizable via a *visual metric*.

Definition 2.17. For $\xi, \chi \in \partial X$ and $x_0 \in X$, we define

$$(\xi \cdot \chi)_{x_0} = \inf_{\substack{(x_i) \in [\xi] \\ (y_i) \in [\chi]}} \liminf_{i, j \rightarrow \infty} (x_i \cdot y_j)_{x_0} \quad (1)$$

Lemma 2.18. *For $\xi, \chi \in \partial X$ and any sequences $(x_i) \in [\xi]$ and $(y_i) \in [\chi]$, we have*

$$\liminf_{i, j \rightarrow \infty} (x_i \cdot y_j)_{x_0} =_{100\delta} (\xi \cdot \chi)_{x_0}.$$

Theorem 2.19. *There exists a metric d_v on ∂X such that there are constants $a > 0$ and $C \geq 1$ depending only on δ so that for all $\xi, \chi \in \partial X$,*

$$\frac{1}{C} a^{-(\xi \cdot \chi)_{x_0}} \leq d_v(\xi, \chi) \leq C a^{-(\xi \cdot \chi)_{x_0}}.$$

The visual metric topology agrees with the topology on ∂X presented earlier.

2.4 Group Actions

We briefly go over the needed assumptions on the action $G \curvearrowright X$. Recall that X is a simplicial graph and is δ -hyperbolic, and that G acts on X by isometries. We take a fixed basepoint $x_0 \in X$, although the definitions given are in fact basepoint-independent.

Definition 2.20. An element $a \in G$ is *loxodromic* if $(u^i x_0)_{i \in \mathbb{Z}}$ is a (λ, c) -quasi-geodesic for some $\lambda \geq 1$ and $c \geq 0$. Two loxodromics $a, b \in G$ are *independent* if there exists some $D > 0$ so that for all $i, j \in \mathbb{Z}$,

$$(u^i x_0 \cdot v^j x_0)_{x_0} < D$$

We now define *acylindrical* and *WPD* (weak proper discontinuity) actions; the definition of acylindricity is due to Osin [Osi16], while the definition of a WPD action is due to Bestvina-Fujiwara [BF02]. The WPD condition is strictly weaker than acylindricity. We will only assume that the actions we study are WPD, but we include the definition of acylindricity as a useful point of comparison.

Definition 2.21. Assume that G acts by isometries on a metric space X . The action $G \curvearrowright X$ is acylindrical if for all $K \geq 0$, there are constants $R(K)$ and $N(K)$ such that for all $x, y \in X$ with $d(x, y) \geq R(K)$, the coarse stabilizer

$$\text{Stab}_K(x, y) := \{g \in G : d_X(x, gx) \leq K \text{ and } d(y, gy) \leq K\}.$$

has cardinality at most $N(K)$.

A specific element $\varphi \in G$ is WPD if for any $x \in X$ and any $K \geq 0$, there exists $N_0 > 0$ such that for $N \geq N_0$, $\text{Stab}_K(x, \varphi^N x)$ is finite.

The action $G \curvearrowright X$ is WPD if G has at least one WPD element.

2.5 Probability Measures, Random Walks, and Axes

We will consider a finite set of probability distributions $\mu_1, \dots, \mu_k$ on G .

Definition 2.22. Given a probability distribution μ on G , (μ, G) is *permissible* if μ has finite support that generates G , the action of G is non-elementary, and the maximal finite normal subgroup of G is trivial. We will usually refer to just a permissible μ on G , with the requirements on G being implied.

Definition 2.23. Given a permissible probability measure μ on G , let us consider the μ -i.i.d. $(g_i)_{i \geq 0} \in G^{\mathbb{Z}_{>0}}$. The *random walk* $(w_n)_{n \geq 0}$ generated by μ is the sequence of random variables defined by

$$w_n := g_1 g_2 \cdots g_n.$$

It is straightforward to check that, for every loxodromic element $g \in G$, the translation length

$$L(g, x) := \lim_{n \rightarrow +\infty} \frac{d(x, g^n x)}{n}$$

is well-defined, is independent of the choice $x \in X$, and is strictly positive. Moreover, when $L(g) \geq 100\delta$, the set

$$\{x \in X : d(x, gx) \leq L(g) + 4\delta\}$$

is nonempty.

Definition 2.24. Let $g \in G$ be a loxodromic element with translation length greater than 100δ . Picking a point x such that $d(x, gx) \leq L(g) + 4\delta$, we define the *axis* $\text{Ax}(g)$ as the concatenation

$$\bigcup_{i=-\infty}^{\infty} [g^i x, g^{i+1} x]_X$$

Definition 2.25. Fix $D > 0$ and $M > 0$, and let $G \curvearrowright X$ be as above. Also fix a parameter n . Given $w_1, \dots, w_k \in G$ satisfying $d_G(e, w_i) \leq n$, we say that $(w_1, \dots, w_k) \in \mathcal{E}_{D,M}(X)$, or $(w_1, \dots, w_k) \in \mathcal{E}_{D,M}$ if the space X is clear from context, if there exist respective conjugates $v_1, \dots, v_k$ of $w_1, \dots, w_k$ satisfying the following properties. In words, we say $(w_1, \dots, w_k)$ are *admissible*.

1. For each i , $(v_i^k x_0)_{k \in \mathbb{Z}}$ is a $(1, \delta)$ -coarse geodesic at speed at least D . In particular, the v_i are loxodromics on X , and each $\text{Ax}(v_i)$ is a 8δ -convex $(1, 4\delta)$ -quasi-geodesic.
2. For each v_i and for each $u \in G$ with $d_G(e, u) \leq M$, there exists an element $g \in G$ such that $(x_0, gx_0, gux_0, v_i x_0)$ is 10δ -aligned, and such that $d_G(e, g)$ and $d_G(e, u^{-1}g^{-1}v_i)$ are bounded above by $2n$.
3. (*Non-overlapping axes.*) For any $i \neq j$, there is no $g \in G$ such that a D/M -long subsegment of $[x_0, v_i x_0]$ and a D/M -long subsegment of $[x_0, v_j x_0]$ are 50δ -fellow travelling.
For any i , there is no $g \in G \setminus \{id\}$ such that a D/M -long subsegment of $[x_0, v_i x_0]$ and a D/M -long subsegment of $g[x_0, v_i x_0]$ are 50δ -fellow traveling.
4. For any v_i , the stabilizer of $\text{Ax}(v_i)$ is $\langle v_i \rangle \times \mathcal{E}(G)$, where $\mathcal{E}(G)$ is the maximal finite normal subgroup of G .

Theorem 2.26. *Let $G, X, \mu_1, \dots, \mu_k$ be as in the setup, and let $(w_n^{(1)})_{n>0}, \dots, (w_n^{(k)})_{n>0}$ be independent random walks generated by $\mu_1, \dots, \mu_k$, respectively. Then there exists $\Delta > 0$ such that for the chosen n and for each $M > 0$,*

$$\lim_{n \rightarrow \infty} \mathbb{P}((w_n^{(1)}, \dots, w_n^{(k)}) \in \mathcal{E}_{\Delta n, M}) = 1.$$

Proof. Among the three items that should be established, the first and the second items are due to Maher and Tiozzo [MT18, Lemma 5.9, Proposition 5.4]. We note that the claim for $\text{Ax}(v_i)$ follows from the coarse geodesic claim by [Cou16], Corollary 2.7 and Lemma 3.2. We note that these items do not require the WPD property of the action, and only requires that the action is non-elementary.

Item 3 follows from combining Proposition 9.12-9.14 of [MT21]. See also Proposition 5.13 of [Abb+26]. Lastly, Item 4 is a direct consequence of Item 3, and it is also recorded in Proposition 5.1 of [MS19].

2.6 Action Path-Connectedness

Lastly, we will need the following assumption. While this assumption follows from the Hierarchy Paths Theorem for hierarchically hyperbolic groups, we state it independently of the full HHG machinery.

Definition 2.27. Let G be a group acting on a geodesic metric space X by isometries, and fix a basepoint x_0 . We say that the action is *action path-connected* with constants E and W if for every pair of points $x, y \in X$, there exists a sequence of group elements $a_0, a_1, \dots, a_k$ such that $a_0 x_0 = x$, $a_k x_0 = y$, such that $d_G(a_{i-1}, a_i) \leq W$ for all $1 \leq i \leq k$, and so that the path $\{a_0 x_0, a_1 x_0, \dots, a_k x_0\}$ is a (E, E) -quasi-geodesic under some reparameterization.

To obtain examples of actions with this property, we apply the following theorem.

Theorem 2.28 (Thm. 6.11, [BHS19]). *Let $\mathcal{X}$ be a hierarchically hyperbolic space. There exists a $\lambda \geq 1$ so that for any $x, y \in \mathcal{X}$, there exists a (λ, λ) -quasi-geodesic γ joining x and y in $\mathcal{X}$ so that for any coordinate space X of the hierarchically hyperbolic space structure for $\mathcal{X}$, the projection $\pi_X \circ \gamma$ is a unparameterized (λ, λ) -quasi-geodesic in X .*

Lemma 2.29. *Let G be a hierarchically hyperbolic group, and let X be the top-level space of a hierarchically hyperbolic space structure realizing G as a hierarchically hyperbolic group. Assume that the action of G on X is transitive. Then the action $G \curvearrowright X$ is action path-connected for some E .*

Proof. For G to be a hierarchically hyperbolic group, the hierarchically hyperbolic space for which X is the top-level space is the Cayley graph $\mathcal{X}$ of G with respect to some finite generating set. Furthermore, we may assume without loss of generality that there is a projection map $\pi_X : \mathcal{X} \rightarrow X$ sending each $a \in \mathcal{X}$ to ax_0 .

Applying Theorem 2.28, there is a constant E so that for any $x, y \in X$, we may find $a_0, a_k \in \mathcal{X}$ such that $a_0x_0 = x$ and $a_kx_0 = y$, and a sequence $a_0, a_1, \dots, a_k$ forming an (E, E) -quasi-geodesic in $\mathcal{X}$ so that its projection $x = a_0x_0, a_1x_0, \dots, a_kx_0 = y$ is a unparameterized (E, E) -quasi-geodesic. As $a_0, a_1, \dots, a_k$ form an E -quasi-geodesic in $\mathcal{X}$, we conclude by the definition of distances in $\mathcal{X}$ that $d_G(a_{i-1}, a_i) \leq W$ for some W dependent only on E . $\square$

3 Hypotheses

In this brief section, we establish our standing hypotheses.

Hypothesis 3.1. *Let G be a finitely generated group acting by isometries on a simplicial graph X with basepoint x_0 , and let $w_1, \dots, w_k \in G$. We assume the following for some fixed constants E and $M > 10^{20}$:*

1. *The space X is Gromov hyperbolic with hyperbolicity constant $\delta > 1000$.*
2. *The action $G \curvearrowright X$ is vertex-transitive, WPD, and non-elementary.*
3. *The action $G \curvearrowright X$ is action path-connected with constant E .*
4. *The maximal finite normal subgroup of G is trivial.*
5. *We have*

$$(w_1, \dots, w_k) \in \mathcal{E}_{10^{20}M, M}$$

as in Definition 2.25.

Given these hypotheses, we set $L = \max_i L(w_i)$ and $\mathcal{M} := 50 \min_i L(w_i)/M$. We also let $v_1, \dots, v_k$ be the respective conjugates of $w_1, \dots, w_k$ obtained from Definition 2.25, and set the notation

$$\mathcal{A} := \{hAx(v) : h \in G, v \in \{v_1, \dots, v_k\}\}.$$

Note that by Theorem 2.26, there exists a constant M for which the last condition of Hypothesis 3.1 holds with probability approaching 1 as n goes to ∞ .

Remark 3.2. We note that several of these hypotheses are more or less free. Specifically, given the other properties and that the action $G \curvearrowright X$ is cobounded, we may assume without loss of generality that $G \curvearrowright X$ is vertex-transitive, that X is a simplicial graph with the simplicial metric, and that the maximal finite normal subgroup of G is trivial, perhaps after adjusting constants. Crucially, these modifications do not affect the conclusions of our theorems. The first two reductions are standard, while the third is explained in Appendix A of [Abb+26].

4 Tight conjugates and minimal representatives

Let $w_1, \dots, w_k \in G$ be elements satisfying Hypothesis 3.1. We will denote

$$L_i = L(w_i) := \lim_n \frac{1}{n} d(x_0, w_i^n x_0).$$

Recall the definition

$$\mathcal{M} := \min_i L(w_i)/M$$

In the following, we will consider the normal closure of w_i 's; in this discussion, we may replace each w_i with any of its conjugates. Hence, without loss of generality, we will assume that the $(w_i^n x_0)_{n \in \mathbb{Z}}$ are 10δ -aligned. What we want is the following.

Lemma 4.1. *Let h be a nontrivial element of $N := \langle\langle w_1, \dots, w_k \rangle\rangle$. Then there exist $g \in G$, $i \in \{1, \dots, k\}$ and a conjugate v of w_i such that $(x_0, gx_0, gv x_0, hx_0)$ is $\mathcal{M}$ -aligned. In particular, for every $x_1, x_2 \in X$ with $hx_1 = x_2$ for $h \in N \setminus \{e\}$, there exists an axis $\gamma \in \mathcal{A}$ such that $d_\gamma(x_1, x_2) \geq L(v) - 3\mathcal{M}$.*

This lemma follows from Lemma 2.1(1) of Clay-Mangahas, but we prove it directly for completeness. Before the proof, let us introduce some terminology.

For $w \in \{w_1^{\pm 1}, \dots, w_k^{\pm 1}\}$, a conjugate $n = gwg^{-1}$ of w is called *tight* if $d_{hAx(v)}(x_0, nx_0) < 0.6L(v)$ for all $v \in \{w_1, \dots, w_k\}$ and $h \in G$, except when $hAx(v) = gAx(w)$.

Lemma 4.2. *Every conjugate n of $w \in \{w_1, \dots, w_k\}$ is a finite product of tight conjugates.*

Proof. We will prove the lemma by induction on $\| \cdot \|_X$, where $\|g\|_X := d_X(x_0, gx_0)$. Let n be a conjugate of an element of $\{w_1, \dots, w_k\}$, and let $n = gwg^{-1}$ be an effective product

decomposition as defined above. Let $h \in G, v \in \{w_1^{\pm 1}, \dots, w_k^{\pm 1}\}$ be an arbitrary pair such that $d_{h\text{Ax}(v)}(x_0, nx_0) > 0.6L(v)$. We can further assume that $\pi_{h\text{Ax}(v)}(x_0)$ appears earlier than $\pi_{h\text{Ax}(v)}(nx_0)$ along $h\text{Ax}(v)$; we can invert the sign of v if not. If $h\text{Ax}(v) = g\text{Ax}(w)$ for all such pairs (h, v) , then n is tight. So let us assume $h\text{Ax}(v) \neq g\text{Ax}(w)$. Let

$$D_1 = d_X(x_0, h\text{Ax}(v)), \quad D_2 = d_X(h\text{Ax}(v), nx_0).$$

The projection of $[x_0, nx_0]$ onto $h\text{Ax}(v)$ has diameter at least $0.6L(v)$, so we can apply Lemma 2.9 and Lemma 2.4. We obtain a subsegment γ of $[x_0, nx_0]$ that lies in the 10δ -neighborhood of $h\text{Ax}(v)$ such that $d(x_0, \gamma) =_{10\delta} D_1$ and $d(nx_0, \gamma) =_{10\delta} D_2$. Note that either $D_1 \leq D_2$ or $D_1 \geq D_2$ holds. Let us focus on the former case, as the latter can be handled in a symmetric way.

Again by applying Lemma 2.9 and Lemma 2.4, we obtain a subsegment γ' of $[x_0, nx_0]$ lying in the 10δ -neighborhood of $g\text{Ax}(w)$. As $v, w \in \mathcal{E}_{10^{20}M, M}$ and $h\text{Ax}(v) \neq g\text{Ax}(w)$, we know that γ and γ' cannot overlap for length $> \mathcal{M}$. This fact plus $D_1 \leq D_2$ implies that

$$D_1 <_{30\delta} \frac{1}{2}(\|n\|_X - L(w)) - L(v) + \mathcal{M}.$$

In this case, the vast majority of γ precedes γ' along $[x_0, nx_0]$, and so we have

$$\text{diam}(\gamma \cap N_{50\delta}([x_0, \pi_{g\text{Ax}(w)}(x_0)])) \geq_{100\delta} l(\gamma) - \mathcal{M} \geq 0.59L(v).$$

In particular, $d_{h\text{Ax}(v)}(x_0, \pi_{g\text{Ax}(w)}(x_0)) \geq 0.58L(v)$. As n stabilizes $g\text{Ax}(w)$, this also implies that

$$d_{nh\text{Ax}(v)}(nx_0, \pi_{g\text{Ax}(w)}(nx_0)) = d_{h\text{Ax}(v)}(x_0, \pi_{g\text{Ax}(w)}(x_0)) \geq 0.58L(v).$$

Now, let us consider the decomposition

$$n = gwg^{-1} = (hvh^{-1}) \cdot (hv^{-1}h^{-1}gwg^{-1}hvh^{-1}) \cdot (hv^{-1}h^{-1}).$$

This decomposition is depicted in Figure 3. First, we have

$$\|hvh^{-1}\| \leq 2D_1 + L(v) < \|n\|_X.$$

Second, we have

$$\begin{aligned} \|hv^{-1}h^{-1}nhvh^{-1}\|_X &\leq_{100\delta} d_X(hv^{-1}h^{-1}x_0, h\text{Ax}(v)) + d_{h\text{Ax}(v)}(hv^{-1}h^{-1}x_0, \pi_{g\text{Ax}(w)}(x_0)) \\ &\quad + d_X(\pi_{g\text{Ax}(w)}(x_0), \pi_{g\text{Ax}(w)}(nx_0)) + d_{nh\text{Ax}(v)}(\pi_{g\text{Ax}(w)}(nx_0), nhv^{-1}h^{-1}x_0) \\ &\quad + d_X(nh\text{Ax}(v), nhv^{-1}h^{-1}x_0) \\ &\leq d_X(x_0, h\text{Ax}(v)) + |d_{h\text{Ax}(v)}(x_0, \pi_{g\text{Ax}(w)}(x_0)) - L(v)| + d_X(\pi_{g\text{Ax}(w)}(x_0), \pi_{g\text{Ax}(w)}(nx_0)) \\ &\quad + |d_{nh\text{Ax}(v)}(\pi_{g\text{Ax}(w)}(nx_0), nx_0) - L(v)| + d_X(nh\text{Ax}(v), nx_0) \\ &\leq_{100\delta} d(x_0, nx_0) - 0.16L(v). \end{aligned}$$

So the three conjugates are of smaller $\|\cdot\|_X$ -norm than n . Induction finishes the proof. $\square$




Figure 3: The inductive tight conjugate decomposition used in the proof of Lemma 4.2. Explicitly, $n = gwg^{-1}$ is decomposed as $(hvh^{-1}) \cdot (hv^{-1}h^{-1}gwg^{-1}hvh^{-1}) \cdot (hv^{-1}h^{-1})$. Due to the projection distance assumption of $d_{hAx(v)}(x_0, \pi_{gAx(w)}(x_0)) \geq 0.58L(v)$, the three geodesics $[x_0, (hvh^{-1})x_0]$, $[(hvh^{-1})x_0, (hvh^{-1})(hv^{-1}h^{-1}gwg^{-1}hvh^{-1})x_0]$, and $[(hvh^{-1})(hv^{-1}h^{-1}gwg^{-1}hvh^{-1})x_0, nx_0]$ each have shorter length than $[x_0, nx_0]$.

Lemma 4.3. *Let $n \in N$ and let m be the minimum number of tight conjugates needed to express n as a product of tight conjugates. Then there exist m tight conjugates*

$$h_1 u_1 h_1^{-1}, \dots, h_m u_m h_m^{-1} \quad (u_i \in \{w_1^\pm, \dots, w_k^\pm\}, h_i \in G)$$

whose product equals n , and a nonempty sequence $1 = l(1) < \dots < l(T)$ such that for

$$\begin{aligned} n_j &:= (h_1 u_1 h_1^{-1}) \cdots (h_j u_j h_j^{-1}) \quad (j = 1, \dots, m), \\ \gamma_i &:= n_{l(i)-1} h_{l(i)} A x(u_{l(i)}) \quad (i = 1, \dots, T), \end{aligned}$$

the following holds.

1. *For points*

$$p_i = \pi_{\gamma_i}(n_{l(i)-1} x_0), \quad q_i = \pi_{\gamma_i}(n_{l(i)} x_0) \quad (i = 1, \dots, T),$$

we have that q_i is $0.5\mathcal{M}$ -close to $[n_{l(i+1)-1} x_0, p_{i+1}]$ for each $i = 1, \dots, T-1$.

2. *We have $d_X(p_i, q_i) \geq L(u_{l(i)})$, $(q_{i-1} \cdot q_i)_{p_i} < 10\delta$ for $i = 1, \dots, T$ (with $q_0 := x_0$) and $(p_i \cdot p_{i+1})_{q_i} \leq 0.7L(u_{l(i)})$ for $i = 1, \dots, T-1$. Moreover, $(p_T \cdot n x_0)_{q_T} < \mathcal{M}$.*

Proof. Let us now prove the lemma by induction. For $m = 1$ we can take the sequence $\{1 = l(1)\}$. Now, suppose that a product n of m tight conjugates satisfies the condition, with conjugates

$$h_1 u_1 h_1^{-1}, \dots, h_m u_m h_m^{-1}$$

and an auxiliary sequence $1 = l(1) < \dots < l(T)$ as in the lemma statement. Now let us consider $n' = h u h^{-1} \cdot n$ for some tight conjugate $h u h^{-1}$. Let $p_0 := \pi_{h A x(u)}(h u^{-1} h^{-1} x_0)$ and $q_0 := \pi_{h A x(u)}(x_0)$.

Case I. Some part of $[q_0, p_0]$ and $[p_i, q_i]$ are 30δ -fellow traveling for length $> 0.1\mathcal{M}$ for some i . Then by the no-overlap condition of Definition 2.25, we have $h A x(u^{-1}) = n_{l(i)-1} h_{l(i)} A x(u_{l(i)})$, with the same orientation. Moreover, we have

$$h u^{-1} h^{-1} = n_{l(i)-1} h_{l(i)} \cdot u_{l(i)} \cdot (n_{l(i)-1} h_{l(i)})^{-1}.$$

This forces that

$$n' = n_{l(i)-1} \cdot (h_{l(i)+1} u_{l(i)+1} h_{l(i)+1}^{-1}) \cdots (h_m u_m h_m^{-1}),$$

and n' becomes a product of at most $m-1$ tight conjugates. Such an element satisfies the conclusion of the lemma by induction.

Case II. The part of $[q_0, p_0]$ that is 20δ -fellow traveling with a subsegment of $[x_0, n x_0]$ has length $< 0.8\mathcal{M}$. This forces that $\pi_{[q_0, p_0]}(n x_0)$ and $\pi_{[q_0, p_0]}(x_0) \subseteq \mathcal{N}_{10\delta}(q_0)$ must have distance at most $0.9\mathcal{M}$, and we get $(p_0 \cdot n x_0)_{q_0} < \mathcal{M}$. Then n' satisfies the conclusion of the lemma with $\{l'(1), l'(2), \dots\} = \{1\}$.

Case III. A part of $[q_0, p_0]$ is 20δ -fellow traveling with a subsegment of $[q_T, x_0]$. Then one can see that $[x_0, q_0]$ contains a subsegment that is 30δ -fellow traveling with $[p_T, q_T]$,

of length $> L(u_t) - 2.9\mathcal{M}$. Namely, we have an axis γ_T that is closer to x_0 than $hAx(u)$ is, whereas $d_{\gamma_T}(x_0, huh^{-1}x_0) > 0.6L(u_T)$. This contradicts the fact that huh^{-1} was a tight conjugate.

Case IV. We are not in Case I, but some part of $[q_0, p_0]$ is 20δ -fellow traveling with a subsegment of $[\pi_{[p_i, q_i]}(p_{i+1}), p_{i+1}]$ for some i , for length $> 0.6\mathcal{M}$.

In this case, we claim that $(p_0 \cdot p_{i+1})_{q_0} < 0.7L(u)$. This will conclude that n' has the desired decomposition with indices $0 < l(i) < \dots < l(T)$.

To see this, remember that $(x_0 \cdot q_i)_{p_i} < 20\delta$ and $(p_i \cdot p_{i+1})_{q_i} \leq 0.7L(u_{l(i)})$. In particular, there are points $P, Q \in [x_0, p_{i+1}]$ with P coming before Q such that $P \in \mathcal{N}_{30\delta}(p_i)$, $Q \in \mathcal{N}_{20\delta}(\pi_{[p_i, q_i]}(p_{i+1}))$ and such that $d(P, Q) > 0.2L(u_i)$.

Now, let $[q', p'] \subseteq [q_0, p_0]$ be the maximal subsegment that is 20δ -fellow traveling with a subsegment of $[\pi_{[p_i, q_i]}(p_{i+1}), p_{i+1}] \subseteq \mathcal{N}_{50\delta}([n_{l(i+1)-1}x_0, n_{l(i+1)}x_0])$. Then we have $d(q', p') < 0.6L(u)$ by the tightness of $h_{l(i)}u_{l(i)}h_{l(i)}^{-1}$. Also, note that $[q_0, q']$ (however short it is) will be 50δ -fellow traveling with part of $[p_i, \pi_{[p_i, q_i]}(p_{i+1})]$. Such fellow traveling should be shorter than $0.1\mathcal{M}$ as Case I is excluded. In short, we have $d(q_0, p') < 0.65L(u)$. It follows that $(p_0 \cdot p_{i+1})_{q_0} < 0.7L(u)$ as desired. Also, note that $d(q_0, [n_{l(i+1)-1}x_0, p_{i+1}]) < d(q_0, q') + 50\delta < 0.5\mathcal{M}$. $\square$

Proof of Lemma 4.1. Given $n \in N$ with the decomposition described in Lemma 4.3, one has that (x_0, p_T, q_T, nx_0) is an $\mathcal{M}$ -aligned sequence with $p_T, q_T \in n_{l(T)-1}h_{l(T)}Ax(u_T)$ and $d(p_T, q_T) = L(u_T)$. $\square$

As projections between axes have $\mathcal{M}$ -bounded diameters by the non-overlapping condition and δ -hyperbolicity, Lemma 4.1 also holds for axes, albeit with slightly worse constants.

Corollary 4.4. *Let $h_1Ax(v_1)$ and $h_2Ax(v_2)$ be distinct axes such that $\pi_{X/N}(h_1Ax(v_1)) = \pi_{X/N}(h_2Ax(v_2))$. Then there exists an axis $hAx(v)$ such that $d_{hAx(v)}(h_1Ax(v_1), h_2Ax(v_2)) \geq L(v) - 10\mathcal{M}$.*

Lemma 4.1 is useful to show that $\overline{X} = X/N$ is hyperbolic, with hyperbolicity constant at most $\max_i(L_i) + 1000\delta$.

Lemma 4.5. *The space X/N is Gromov hyperbolic, with hyperbolicity constant at most $\max_i(L_i) + 1000\delta$.*

Proof. By Papasoglu's theorem [Pap95, Theorem 1.4], it suffices to prove that all geodesic bigons are δ -coarsely $\max_i(L_i)$ -thin. In fact, we will prove that:

Let $\bar{x}, \bar{y}, \bar{z} \in X/N$ be such that $d_{X/N}(\bar{y}, \bar{z}) < 100\delta$. Then $[\bar{x}, \bar{y}]$ and $[\bar{x}, \bar{z}]$ are within Hausdorff distance $\leq \max_i L_i + 1000\delta$.

To show this, let us lift $[\bar{x}, \bar{y}]$, $[\bar{y}, \bar{z}]$ and $[\bar{z}, \bar{x}]$ to X . More precisely, denoting the projection from X to X/N by π_N , let $[\bar{x}, \bar{y}] = \pi_N(x = x_0, x_1, \dots, x_n)$ for some X -path $(x_0, \dots, x_n)$.




Figure 4: The decomposition defined in Lemma 4.3. In the case depicted, $T = 2$, and we have extended the construction by defining $n' = (huh^{-1})n$. As $[q_0, p_0]$ fellow-travels a segment of $[\pi_{[p_1, q_1]}(p_2), p_2]$, we are in Case IV. We therefore adjust the auxiliary sequence to $0 = l(0) < l(2)$. In particular, we have that $\pi_{\gamma_2}((hu^{-1}h^{-1})x_0, nx_0) \geq L(u_2) - 3\mathcal{M}$.

Similarly, pick $(x_n, \dots, x_{n'})$ and $(x_{n'}, \dots, x_{n''})$ that respectively project to $[\bar{y}, \bar{z}]$ and $[\bar{z}, \bar{x}]$. Then x_0 and $x_{n''}$ are identified under projection, so there exists $n \in N$ such that $nx_0 = x_{n''}$. Here, because $[\bar{x}, \bar{y}]$ is a geodesic, $(x_0, \dots, x_n)$ is also a geodesic from x_0 to x_n . Similarly, $(x_n, \dots, x'_n), (x_{n'}, \dots, x_{n''})$ are also geodesics.

Now, by Lemma 4.1, there exists $g \in G$ and a conjugate v of $\{w_1^{\pm 1}, \dots, w_k^{\pm 1}\}$ such that (x_0, gx_0, gvx_0, nx_0) is $L(v)/50$ -aligned. Let $\gamma \subseteq [x_0, nx_0]$ be the subsegment of $[x_0, nx_0]$ contained in the 50δ -neighborhood of $[gx_0, gvx_0]$. If $N_{50\delta}(\gamma)$ intersects the geodesic $(x_0, \dots, x_n)$ for length $> 0.6L(v)$, then $d_X(gvg^{-1}x_0, x_n)$ becomes shorter than $d_X(x_0, x_n)$. This gives a contradiction, as then $d_{X/N}(\bar{x}, \bar{y}) = d_X(x_0, x_n) < d_X(gvg^{-1}x_0, x_n) = d_{X/N}(\bar{x}, \bar{y})$. (Note that $\pi(x_0) = \pi(gvg^{-1}x_0)$.) Likewise, $N_{50\delta}(\gamma)$ cannot intersect the geodesic $(x_{n'}, \dots, x_{n''})$ for length more than $0.6L(v)$.

It follows that the beginning point of γ is 50δ -close to a point x_i on $(x_0, \dots, x_n)$ and the ending point of γ is 50δ -close to a point x_j on $(x_{n'}, \dots, x_{n''})$. Clearly, $\pi(x_i)$ and $\pi(x_j) = \pi(gv^{-1}g^{-1}x_i)$ are 100δ -close. Now, for the quadrangle with vertices $x_i, x_n, x_{n'}, x_j$, we know that $d(x_n, x_{n'}) \leq 100\delta$ and $d(x_i, x_j) < L(v) + 100\delta$. For such quadrangles, $[x_i, x_n]$ and $[x_{n'}, x_j]$ are $L(v) + 1000\delta$ -close in the Hausdorff distance (even at the level of X). Now we can apply our inductive hypothesis to $y' = \pi(x_i)$ and $z' = \pi(x_j)$. $\square$

5 Strut Sequences

For this section, it will be important to sometimes replace projections with the modified projections defined by Bestvina–Bromberg–Fujiwara–Sisto in [Bes+19], and measure projection distances with respect to these modified projections. These modified projection distances capture a ‘stability’ property of closest point projections with respect to translation along axes, and thus make certain inductive arguments cleaner.

Recall the following definition:

Definition 5.1. Given $w_1, \dots, w_k$ as in Hypothesis 3.1, let $v_1, \dots, v_k$ be as in Definition 2.25. We define the set $\mathcal{A}$ of *axes* by

$$\mathcal{A} = \{hAx(v) : h \in G, v \in \{v_1, \dots, v_k\}\}.$$

Note that $hAx(v) = Ax(hvh^{-1})$.

Note that distinct axes in $\mathcal{A}$ have mutual projections with diameter $\leq 0.02\mathcal{M}$, by the 5th condition in Hypothesis 3.1. This also gives the Behrstock inequality (given that $L(v_i)$ ’s are all long enough so that $\mathcal{M} > 1000\delta$):

$$d_\gamma(\gamma', \gamma'') \geq 0.05\mathcal{M} \Rightarrow d_{\gamma'}(\gamma, \gamma'') \leq 0.05\mathcal{M} \quad (\forall \gamma, \gamma', \gamma'' \in \mathcal{A}).$$

These facts lead to the following consequence.

Lemma 5.2 ([Bes+19, Theorem 4.1]). *There exist modified projection functions Π and modified distance functions D such that, for any axes α and β and $x, y, z \in X \cup \mathcal{A}$:*

1. We have $\Pi_\alpha(x) \subseteq \mathcal{N}_\mathcal{M}(\pi_\alpha(x))$ and $\text{diam}(\Pi_\alpha(x)) \leq \mathcal{M}$.

2. We have

$$D_\alpha(x, y) =_\theta d_\alpha(x, z)$$

3. We have

$$D_\alpha(x, y) = \text{diam}(\Pi_\alpha(x) \cup \Pi_\alpha(y))$$

4. For any $g \in G$, $\Pi_{g\alpha}(gx) = g\Pi_\alpha(x)$, so as a consequence $D_\alpha(gx, gy) = D_\alpha(x, y)$.

5. We have $D_\alpha(x, z) \leq D_\alpha(x, y) + D_\alpha(y, z)$.

6. (Behrstock Inequality) We have

$$\min(D_\alpha(x, \beta), D_\beta(x, \alpha)) \leq \mathcal{M}.$$

7. The set $\{\eta \in \mathcal{A} : D_\eta(x, y) \geq \mathcal{M}\}$ is finite.

8. (Strong Behrstock Inequality / ‘Forcing’) If $d_\beta(x, \alpha) > \theta$, then $\Pi_\alpha(x) = \Pi_\alpha(\beta)$. As a consequence,

$$d_\beta(x, \alpha) > \mathcal{M} \text{ implies } D_\alpha(x, w) = D_\alpha(\beta, w)$$

This last item is generally called the Strong Behrstock Inequality, although we will use the term ‘forcing’ to describe its conclusion.

Remark 5.3. From here onwards, we increase δ so that for all $\alpha \in \mathcal{A}$ and $x \in X \cup \mathcal{A}$, $\mathcal{N}_{10\delta}(\pi_\alpha(x)) \subseteq \Pi_\alpha(x)$.

We would like to be able to lift geodesic triangles from X/N to X ; unfortunately, this is usually not possible. Instead, we will lift triangles that are *near-geodesic* in the following sense.

Definition 5.4. Let $x, y \in X \cup \mathcal{A}$, with projections $\bar{x} = \pi_{X/N}(x)$ and $\bar{y} = \pi_{X/N}(y)$. We say that $\bar{x}$ and $\bar{y}$ are *minimal* if $d_X(x, y) = d_{X/N}(\bar{x}, \bar{y})$. Furthermore, we say that x and y are τ -near-minimal for $0 < \tau < 0.99$ if for all axes $hAx(v)$, $D_{hAx(v)}(x, y) \leq \tau L(v)$. If no τ is stated, we assume $\tau = 0.8$.

Lemma 5.5. If x and y are minimal, then for all axes $hAx(v)$, $d_{hAx(v)}(x, y) \leq 0.51L(v)$, so in particular x and y are near-minimal. Furthermore, if x and y are τ -near-minimal, then for all axes $hAx(v)$, $d_{hAx(v)}(x, y) \leq (\tau + 0.01)L(v)$.

Proof. For any axis $hAx(v)$, there exists some integer m so that $d_{hAx(v)}(x, gy) \leq 0.51L(v)$ for $g = hv^mh^{-1}$. In fact, this upper bound may be sharpened to $0.505L(v)$. If x and y are minimal but $d_{hAx(v)}(x, y) > 0.51L(v)$ for some axis $hAx(v)$, then it is straightforward to check from the Bounded Geodesic Image Theorem that $d_X(x, gy) < d_X(x, y)$, contradicting minimality. The second claim is immediate from Lemma 5.2. $\square$

Lemma 5.6. Fix $\bar{x}, \bar{y} \in X/N \cup \mathcal{A}/N$, along with respective lifts x and y in $X \cup \mathcal{A}$. Assume that for all axes $hAx(v)$, $d_{hAx(v)}(x, y) \leq 0.45L(v)$. Then x and y are minimal.

Proof. Let x' be a minimal lift of $\bar{x}$ with respect to y . If $x = x'$, we're done. Otherwise, by Lemma 4.1, there exists an axis $hAx(v)$ such that $d_{hAx(v)}(x, x') \geq 0.99L(v)$. We have $d_{hAx(v)}(x, y) \leq 0.45L(v)$ by assumption, and $d_{hAx(v)}(x', y) \leq 0.51L(v)$ by minimality, so by the triangle inequality $d_{hAx(v)}(x, x') \leq 0.96L(v)$. This is a contradiction. $\square$




Figure 5: A representative illustration of Lemma 5.10. The left pane shows the quotient space X/N , while the right pane shows X . While lifting the geodesic triangle between $\bar{x}$, $\bar{y}$, and $\bar{z}$ from X/N to X is not possible, we may ‘close it up’ by successive ‘rotations’ along axes. In the quotient, these rotations correspond to ‘rerouting’ paths around axes.

We will now define *strut sequences*, which will be a primary technical tool for the rest of the paper. This machinery captures the necessary small-cancellation properties of axes that we need.

Definition 5.7 (Strut Sequences). Let $x, y \in X \cup \mathcal{A}$, and assume x and y are near-minimal. For $0 < \theta < 1$, the θ -strut sequence for x and y is

$$\mathcal{S}_\theta(x, y) = \{hAx(v) \in \mathcal{A} : D_{hAx(v)}(x, y) \geq \theta L(v)\} \cup \{x, y\}.$$

We define a finite total order $<_{\mathcal{S}}$ on $\mathcal{S}_\theta(x, y)$ by defining x to be the minimal element and then setting $\alpha <_{\mathcal{S}} \beta$ if $d_\alpha(x, \beta) > 10\mathcal{M}$. If no θ is stated, then it is implied that $\theta = 0.1$.

For an illustration of strut sequences in X and in X/N , see Figure 6. Note that the visual representation in this figure assumes several facts that have not yet been proven.

We now set up some basic facts about strut sequences.

Lemma 5.8. *We have the following properties for any $x, y \in X \cup \mathcal{A}$ and any $0.01 \leq \theta \leq 0.4$. Assume that $\mathcal{S}_\theta(x, y)$ is enumerated as*

$$x <_{\mathcal{S}} h_1Ax(v_1) <_{\mathcal{S}} \dots <_{\mathcal{S}} h_tAx(v_t) <_{\mathcal{S}} y$$

1. *For any $x, y \in X \cup \mathcal{A}$, the total order $<_{\mathcal{S}}$ defined in Definition 5.7 is well-defined, and $\mathcal{S}_\theta(x, y)$ is finite.*
2. *For any $h_iAx(v_i)$ and any $j_1 < i < j_2$,*

$$\Pi_{h_iAx(v_i)}(x) = \Pi_{h_iAx(v_i)}(h_{j_1}Ax(v_{j_1})), \quad \Pi_{h_iAx(v_i)}(h_{j_2}Ax(v_{j_2})) = \Pi_{h_iAx(v_i)}(y)$$

3. *We have that $<_{\mathcal{S}}$ -consecutive elements are minimal. If x and y are near-minimal, then any two elements in $\mathcal{S}_\theta(x, y)$ are near-minimal.*
4. *If $w, z \in \mathcal{S}_\theta(x, y)$ with $w <_{\mathcal{S}} z$, then $\mathcal{S}_\theta(w, z)$ is equal to the suborder of $\mathcal{S}_\theta(x, y)$ bounded by w and z .*

Proof. Finiteness follows immediately from Lemma 5.2. Now, for any $hAx(v) \in \mathcal{S}_\theta(x, y)$, there must be some maximal segment $\gamma_{hAx(v)}$ of γ such that $\gamma_{hAx(v)}$ 10δ -fellow-travels $hAx(v)$. By the Bounded Geodesic Image Theorem, these segments each have length at least $0.09L(v)$, and by the no-overlap hypothesis for axes (item 3 in Definition 2.25), may only overlap for distance $10\mathcal{M}$. Their order on γ gives a finite total order, and items 6 and 8 of Lemma 5.2 implies this order is equal to the order $<_{\mathcal{S}}$.

The second item now follows immediately from the forcing property in Lemma 5.2. The proofs of the third and fourth claim are omitted. $\square$

The following useful lemma follows near-immediately from the definitions and the forwarding property for projections.

Lemma 5.9. *Fix $0.01 \leq \theta \leq 0.4$, $0.55 \leq \tau \leq 0.95$, $x, y \in X \cup \mathcal{A}$ and $hAx(v) \in \mathcal{A}$ such that $\theta \leq D_{hAx(v)}(x, y)/L(v) \leq \tau$. Then $\mathcal{S}_\theta(x, y)$ is obtained by concatenating $\mathcal{S}_\theta(x, hAx(v))$, $\{hAx(v)\}$, and $\mathcal{S}_\theta(hAx(v), y)$ in that order. Furthermore, if $hAx(v)$ is near-minimal with respect to both x and y , then x and y are near-minimal.*




Figure 6: A depiction of strut sequences in X and under projection to X/N , and how these sequences control routing of geodesics between near-minimal pairs. The lower pane shows a point $x \in X$ and two points y_1, y_2 in a common orbit, such that y_1 and x are minimal and y_2 and x are near-minimal. In X , the strut sequences $\mathcal{S}(x, y_i)$ represent sequences of axes on which the $[x, y_i]$ make ‘large turns’. In X/N , we see that the paths project to two paths between $\bar{x}$ and $\bar{y}$, which 50δ -fellow-travel except when routing around the quotient axes obtained by projecting either strut sequence. A heuristic to keep in mind is that while X/N is **not** δ -hyperbolic, the failure of δ -hyperbolicity only comes from looping around quotient axes. In particular, the two geodesics $\bar{\gamma}_1$ and $\bar{\gamma}_2$ must approximately δ -fellow-travel between the struts, allowing us to recover some features of δ -hyperbolicity in X/N .

The primary reason to study near-minimal representatives is the following lemma. This result is conceptually analogous to Prop. 3.12 in [Abb+26].

Lemma 5.10. *Let $\bar{x}, \bar{y}, \bar{z} \in X/N$. Also let $0.78 \leq \tau < 0.95$. Then there exist respective lifts $x, y, z \in X$ of $\bar{x}, \bar{y}$, and $\bar{z}$ such that y and z are minimal, x and y are τ -near-minimal, and x and z are τ -near-minimal.*

Proof. Choose lifts x_1, y, z , and x_2 so that x_1 and y, y and z , and z and x_2 are all minimal. We induct on $d_X(x_1, x_2)$, maintaining at each step that y and x_1 are τ -near-minimal and that z and x_2 are τ -near-minimal. If $d_X(x_1, x_2)$ is ever equal to 0, we finish immediately.

Otherwise, by Lemma 4.1, there exists an axis $hAx(v)$ such that $d_{hAx(v)}(x_1, x_2) \geq 0.99L(v)$. We consider the projections of x_1, y, z , and x_2 to $hAx(v)$. By assumption, we have $d_{hAx(v)}(x_1, y) \leq (\tau + 0.01)L(v)$, $d_{hAx(v)}(z, x_2) \leq (\tau + 0.01)L(v)$, and $d_{hAx(v)}(y, z) \leq 0.51L(v)$. We may also assume that v is oriented so that the projection order on $hAx(v)$ is $x_1 \rightarrow x_2$.

From the distance bounds, either $d_{hAx(v)}(x_1, y) \geq 0.24L(v)$ and we have the projection order $x_1 \rightarrow y \rightarrow x_2$, or $d_{hAx(v)}(z, x_2) \geq 0.24L(v)$ and we have the projection order $x_1 \rightarrow z \rightarrow x_2$. Assume without loss of generality that we are in the first case. Then if $g = hvh^{-1}$, by the distance bounds and the translation length of g we have

$$(0.99 - \tau)L(v) \leq d_{hAx(v)}(gx_1, y) \leq 0.77L(v).$$

By the inductive hypothesis plus Lemma 5.8, we have that gx_1 and $ghAx(v) = hAx(v)$ are τ -near-minimal, and that y and $hAx(v)$ are τ -near-minimal. We may therefore apply Lemma 5.9 and conclude gx_1 and y are near-minimal. Furthermore, by the Bounded Geodesic Image Theorem, $d_X(gx_1, x_2) < d_X(x_1, x_2)$, so we may apply induction and conclude. $\square$

This ‘rotation’ action is depicted in Figure 5, although with g applied to x_2 .

We may now prove a crucial result about how struts sequences behave under quotients. See Figure 6 for an illustration of this property.

Lemma 5.11. *Fix $0.55 < \tau < 0.95$ and $0.01 < \theta < 0.4$, and assume $\theta + \tau < 0.98$. Fix $\bar{x}, \bar{y} \in X/N$, and consider lifts $x_1, x_2 \in \pi_{X/N}^{-1}(\bar{x})$ and $y_1, y_2 \in \pi_{X/N}^{-1}(\bar{y})$. Assume that x_1 and y_1 are τ -near-minimal, and that x_2 and y_2 are τ -near-minimal. Enumerate strut sequences*

$$\mathcal{S}_\theta(x_1, y_1) = x_1 <_{\mathcal{S}} h_1 Ax(v_1) <_{\mathcal{S}} \dots <_{\mathcal{S}} h_t Ax(v_t) <_{\mathcal{S}} y_1$$

$$\mathcal{S}_\theta(x_2, y_2) = x_2 <_{\mathcal{S}} h'_1 Ax(v'_1) <_{\mathcal{S}} \dots <_{\mathcal{S}} h'_{t'} Ax(v'_{t'}) <_{\mathcal{S}} y_2$$

Then $t = t'$ and $\pi_{X/N}(h_i Ax(v_i)) = \pi_{X/N}(h'_i Ax(v'_i))$ for all $h_i Ax(v_i)$.

Furthermore, fix $i < j$, and consider $h_i \text{Ax}(v_i), h_j \text{Ax}(v_j), h'_i \text{Ax}(v'_i), h'_j \text{Ax}(v'_j)$ with the assumption that for all $i < k < j$,

$$D_{h_k \text{Ax}(v_k)}(h_i \text{Ax}(v_i), h_j \text{Ax}(v_j)) \leq 0.48L(v_k).$$

Then there is a single $h \in N$ such that

$$h\pi_{h_i \text{Ax}(v_i)}(h_{i+1} \text{Ax}(v_{i+1})) = \pi_{h'_i \text{Ax}(v'_i)}(h'_{i+1} \text{Ax}(v'_{i+1}))$$

and

$$h\pi_{h_{i+1} \text{Ax}(v_{i+1})}(h_i \text{Ax}(v_i)) = \pi_{h'_{i+1} \text{Ax}(v'_{i+1})}(h'_i \text{Ax}(v'_i))$$

Proof. We consider the first claim, assuming without loss of generality that $y_1 = y_2 =: y$. The strategy is induction on $\max(d_X(x_1, y_1), d_X(x_2, y_2))$; if both are less than, say, $0.01 \min_v L(v)$, then by Lemma 4.1 and the triangle inequality we must have $x_1 = x_2$, and are done. If $x_1 \neq x_2$, by Lemma 4.1 there exists an axis $h \text{Ax}(v)$ such that $d_{h \text{Ax}(v)}(x_1, x_2) \geq 0.99L(v)$. By assumption, $d_{h \text{Ax}(v)}(x_i, y) \leq \tau$ for $i \in \{1, 2\}$, so by the triangle inequality $d_{h \text{Ax}(v)}(x_i, y) \geq 0.98 - \tau$. Therefore $h \text{Ax}(v) \in \mathcal{S}_\theta(x_i, y)$. Furthermore, certainly $d_X(x_i, h \text{Ax}(v)) < d_X(x_i, y)$, so we have the claim for $\mathcal{S}_\theta(x_i, h \text{Ax}(v))$. Applying Lemma 5.9 with $h \text{Ax}(v)$, $\mathcal{S}_\theta(x_i, h \text{Ax}(v))$, and $\mathcal{S}_\theta(h \text{Ax}(v), y)$ finishes the argument.

For the second claim, we may assume without loss of generality that $h_i \text{Ax}(v_i) = h'_i \text{Ax}(v'_i)$. Certainly there exists some $n = h_i v_i^m h_i^{-1}$ such that $d_{h_i \text{Ax}(v_i)}(nh_j \text{Ax}(v_j), h'_j \text{Ax}(v'_j)) \leq 0.52L(v_i)$. The assumed bounds and the triangle inequality now imply that for all $h \text{Ax}(v)$, $d_{h \text{Ax}(v)}(nh_j \text{Ax}(v_j), h'_j \text{Ax}(v'_j)) \leq 0.98L(v)$, and so by Lemma 4.1 we have $nh_j \text{Ax}(v_j) = h'_j \text{Ax}(v'_j)$. The claim follows. $\square$

6 Routing Geodesics in the Quotient

We now establish consequences of the results we have obtained so far. The following claim is illustrated in Figure 6.

Lemma 6.1. Fix $0.55 < \tau < 0.95$ and $0.01 < \theta < 0.4$, and assume $\theta + \tau < 0.98$. Fix $\bar{x}, \bar{y} \in X/N$, and consider lifts $x_1, x_2 \in \pi_{X/N}^{-1}(\bar{x})$ and $y_1, y_2 \in \pi_{X/N}^{-1}(\bar{y})$. Assume that x_1 and y_1 are τ -near-minimal, and that x_2 and y_2 are τ -near-minimal, and choose geodesics $[x_1, y_1]$ and $[x_2, y_2]$. Then the Hausdorff distance of $\pi_{X/N}([x_1, y_1])$ and $\pi_{X/N}([x_2, y_2])$ is at most $\max_v L(v) + 1000\mathcal{M}$. Furthermore,

$$d(x_1, y_1) \leq 1.01 \frac{\tau}{\theta} d(x_2, y_2)$$

Proof. Enumerate struts as follows:

$$\mathcal{S}_\theta(x_1, y_1) = x_1 <_{\mathcal{S}} h_1 \text{Ax}(v_1) <_{\mathcal{S}} \dots <_{\mathcal{S}} h_t \text{Ax}(v_t) <_{\mathcal{S}} y_1$$

$$\mathcal{S}_\theta(x_2, y_2) = x_2 <_S h'_1 \text{Ax}(v'_1) <_S \dots <_S h'_t \text{Ax}(v'_t) <_S y$$

By the Bounded Geodesic Image Theorem, the geodesic $[x_1, y_1]$ may be decomposed into segments γ_i that each 10δ -fellow-travel $h_i \text{Ax}(v_i)$ and interstitial segments η_i that each originate in $\mathcal{N}_{10\delta}(\pi_{h_i \text{Ax}(v_i)}(h_{i+1} \text{Ax}(v_{i+1})))$ and terminate in $\mathcal{N}_{10\delta}(\pi_{h_{i+1} \text{Ax}(v_{i+1})}(h_i \text{Ax}(v_i)))$. We construct analogous segments γ'_i and η'_i for $[x_2, y_2]$. Note that this notation breaks slightly at the endpoints; these cases follow from essentially the same argument.

We now apply Lemma 5.11. Under projection to X/N , any γ_i and γ'_i both lie in $\mathcal{N}_{10\delta}(h_i \text{Ax}(v_i))$; this set has diameter at most $L(v_i) + 100\delta$. Furthermore, both γ_i and γ'_i have length at least $\theta - 100\mathcal{M}$ and at most $\tau + 100\mathcal{M}$.

From the projection complex bounds and the second claim of Lemma 5.11, we obtain that any η_i and η'_i $10\mathcal{M}$ -fellow-travel, and have lengths that differ by no more than $10\mathcal{M}$.

We may now assemble these results to obtain the claims. If in fact there are no struts between x_1 and y_1 , then $[x_1, y_1]$ and $[x_2, y_2]$ in fact 10δ -fellow-travel after projection to X/N . $\square$

For our main construction, we will need to build our own strut sequences, and obtain fairly precise control over those sequences. We do so as follows.

Lemma 6.2. *Fix $\bar{x}, \bar{y} \in X/N$. Given any fixed lift $x \in \pi_N^{-1}(\bar{x})$, there exist two (possibly non-distinct) points $p_1 = p_1(x, \bar{y})$, $p_2 = p_2(x, \bar{y})$ in X with the following properties:*

1. *The distances $d_X(x, p_1)$ and $d_X(x, p_2)$ are bounded above by $\frac{\mathcal{M}}{10} + 1$ and bounded below by $\min(\frac{\mathcal{M}}{10} - 1, d_X(x, y))$.*
2. *For any $y \in \pi_N^{-1}(\bar{y})$ such that x and y are near-minimal, either $(x \cdot y)_{p_1} \leq 100\delta$ or $(x \cdot y)_{p_2} \leq 100\delta$*

Proof. Choose $y \in \pi_{X/N}^{-1}(\bar{y})$ so that x and y are minimal. Enumerate the strut sequence for x and y as

$$\mathcal{S}(x, y) = x <_S h_1 \text{Ax}(v_1) <_S \dots <_S h_t \text{Ax}(v_t) <_S y$$

If there are no struts between x and y (i.e. $t = 0$), then we choose $p_1 = p_2$ to be a point on $[x, y]$ satisfying the distance bounds. Otherwise, we take a geodesic $[x, q]$ to a point q on $h_1 \text{Ax}(v_1)$ realizing $d_X(x, h_1 \text{Ax}(v_1))$, and then ‘branch’ this geodesic along $h_1 \text{Ax}(v_1)$ in both possible directions for distance $0.01L(v_1)$ to obtain points q_1 and q_2 . We then choose p_1 and p_2 to be points along $[x, q_1]$ and $[x, q_2]$ satisfying the appropriate length conditions.

For any $y' \in \pi_{X/N}^{-1}(\bar{y})$ such that x and y' are near-minimal, then the last claim of Lemma 5.8 implies that both $[x, y]$ and $[x, y']$ have $h_1 \text{Ax}(v_1)$ as a initial strut. So $[x, y']$ must pass within 10δ of $\pi_{h_1 \text{Ax}(v_1)}(x)$ and then follow $h_1 \text{Ax}(v_1)$ for at least distance $0.09L(v_1)$ in one of the two possible directions, and the claim follows. $\square$

Definition 6.3. Given $\bar{x}$, $\bar{y}$, and $x \in \pi_N^{-1}(\bar{x})$ as in Lemma 6.2, we call the $p_1 = p_1(x, \bar{y})$ and $p_2 = p_2(x, \bar{y})$ constructed in this lemma *enemy prongs* for x relative to $\bar{y}$.

Corollary 6.4. Fix $\bar{x}, \bar{y} \in X/N$, choose a fixed lift $x \in \pi_N^{-1}(\bar{x})$, and construct enemy prongs p_1 and p_2 as in Lemma 6.2. Assume that $d_{X/N}(\bar{x}, \bar{y}) \geq 1000L$, and let $\bar{z} \in X/N$ be a point satisfying $d_{X/N}(\bar{y}, \bar{z}) \leq 200L$. Let y and z be respective lifts of $\bar{y}$ and $\bar{z}$ such that x and y are near-minimal and y and z are near-minimal. Then either $(x \cdot z)_{p_1} \leq 110\delta$ or $(x \cdot z)_{p_2} \leq 110\delta$.

Proof. By the assumed distance bounds and Lemma 6.1, $d_X(x, y) \geq 1000L$ while $d_X(y, z) \leq 200L$. Furthermore, by Lemma 6.2, either $(x \cdot y)_{p_1} \leq 200\delta$ or $(x \cdot y)_{p_2} \leq 1000\delta$, and $d_X(x, p_1)$, $d_X(x, p_2)$ are bounded above by, say, L . The claim now follows from a standard δ -thinness argument applied to a triangle with vertices x , y , and z . $\square$

We apply the enemy prongs construction as follows. See Figure 7 for an illustration of the following lemma.

Lemma 6.5. Let B be a large enough constant determined only in terms of the action $G \curvearrowright X$. Let n be a sufficiently large number, let $w_n^{(1)}, \dots, w_n^{(k)}$ be the n -step elements of independent permissible random walks on G and let $N = \langle\langle w_n^{(1)}, \dots, w_n^{(k)} \rangle\rangle$.

Fix $\bar{x}, \bar{y}, \bar{z} \in X/N$ and fix respective lifts $x \in \pi_N^{-1}(\bar{x})$ and $z \in \pi_N^{-1}(\bar{z})$, and let p_1, p_2 be constructed from x and $\bar{y}$ as in Lemma 6.2. Make the following assumptions:

1. We have $(p_1 \cdot z)_x \leq B$, $(p_2 \cdot z)_x \leq B$.
2. We have that x and z are near-minimal.
3. Further assume that any axis $hAx(v)$ that 50δ -fellow-travels $[x, z]_X$ for at least $0.0001L(v)$ does not $10B$ -fellow-travel $[x, p_1]_X$ or $[x, p_2]_X$ for distance more than $100B$.

Then $[\bar{z}, \bar{y}]$ passes within $1000B$ of $\bar{x}$, and in fact for any near-minimal respective lifts z' of $\bar{z}$ and y' of $\bar{y}$ and any $[z', y']$, this geodesic passes within $1000B$ of some lift x' of $\bar{x}$.

Furthermore, if $d_{X/N}(\bar{x}, \bar{y}) > 1000B$ and $\bar{w}$ is some point within distance $200B$ of $\bar{y}$, then $[\bar{z}, \bar{w}]$ passes within $2000B$ of $\bar{x}$.

Proof. We note that Figure 7 illustrates many components of the construction used in this proof.

Let y be a minimal lift of $\bar{y}$ with respect to x . We then label the strut sequences

$$\begin{aligned} \mathcal{S}(z, x) &:= \{x <_{\mathcal{S}} h_1 Ax(v_1) <_{\mathcal{S}} \dots <_{\mathcal{S}} h_s Ax(v_s) <_{\mathcal{S}} x\}, \\ \mathcal{S}(x, y) &:= \{x <_{\mathcal{S}} j_1 Ax(w_1) <_{\mathcal{S}} \dots <_{\mathcal{S}} j_t Ax(w_t) <_{\mathcal{S}} y\}. \end{aligned}$$

We may assume without loss of generality that $d_{X/N}(\bar{x}, \bar{y}) > 500B$, and so p_1 and p_2 are both at least distance $400B$ from x . Also assume that n is sufficiently large so that

$\min_i L(v_i) \geq 10^6 B$. By Lemma 6.2, as x and y are near-minimal, we have that either $(x \cdot y)_{p_1} \leq 100\delta$ or $(x \cdot y)_{p_2} \leq 100\delta$; assume we are in the first case. Recall that by assumption we have $(p_1 \cdot z)_x \leq B$. Taking arbitrary geodesics $[z, x]$, $[x, y]$, and $[z, y]$, we have that p_1 lies within 100δ of $[x, y]$, and that x lies within B of $[z, p_1]$. Recall that $d_X(x, p_1) \geq 400B$. Applying the second and third properties of Lemma 2.11, we obtain that $(z \cdot y)_x \leq 2B$ and $(z \cdot y)_{p_1} \leq 150\delta$.

Applying hyperbolicity, we obtain that $\gamma := [z, y]$ may be subdivided into an arc γ_z that 10δ -fellow-travels $[z, x]$ outside of $\mathcal{N}_{3B}(\{x\})$, and a disjoint arc γ_y that 10δ -fellow-travels $[x, y]$ outside of $\mathcal{N}_{3B}(\{x\})$. Furthermore, γ_y passes within 200δ of p_1 .

We note that any axis $h\text{Ax}(v)$ that 50δ -fellow-travels γ_z for distance at least $0.0002L(v)$ must 50δ -fellow-travel $[x, z]_X$ for distance at least $0.0001L(v)$. But then such a geodesic can only 50δ -fellow-travel γ_y for distance at most $105B$. As a consequence, any strut $h\text{Ax}(v) \in \mathcal{S}(x, z)$ satisfies $D_{h\text{Ax}(v)}(x, y) \leq 0.01L(v)$, and vice-versa. As a consequence, z and y must be 0.81 -near-minimal.

Next, from properties of γ_z and γ_y and a axis-fellow-traveling argument, we see that for any $h_i\text{Ax}(v_i)$, $D_{h_i\text{Ax}(v_i)}(z, y) \geq 0.099L(v_i)$, while for any $j_i\text{Ax}(w_i)$, $D_{j_i\text{Ax}(w_i)}(z, y) \geq 0.099L(w)$. So every strut in $\mathcal{S}(z, x)$ and $\mathcal{S}(x, y)$ lies in $\mathcal{S}_{0.099}(z, y)$. Furthermore, as γ_z precedes γ_x , we must have $h_s\text{Ax}(v_s) <_S j_1\text{Ax}(w_1)$, and as a consequence the strut sequence $\mathcal{S}_{0.099}(z, y)$ contains (as a subsequence) all struts in $\mathcal{S}(z, x)$ followed by all struts of $\mathcal{S}(x, y)$. By Lemma 5.11, we have the second claim. If any strut $h\text{Ax}(v)$ lies between $h_s\text{Ax}(v_s)$ and $j_1\text{Ax}(w_1)$, then as it did not appear in $\mathcal{S}(z, x)$ or $\mathcal{S}(x, y)$, and cannot 10δ -fellow-travel both γ_z and γ_y for distance at least $0.01L(v)$, we must have that $D_{h\text{Ax}(v)}(h_s\text{Ax}(v_s), j_1\text{Ax}(w_1)) \leq 0.12L(v)$.

Now, recall the procedure for constructing p_1 and p_2 : we construct a geodesic from x to $j_1\text{Ax}(w_1)$, branch this geodesic in both directions along $j_1\text{Ax}(w_1)$ to produce two geodesics, and then take p_1 and p_2 as points along this geodesic with (in this case) distance at least $400B$ from x .

Assume now that $h_s\text{Ax}(v_s)$ and $j_1\text{Ax}(w_1)$ 20δ -fellow-travel for distance $400B$, and let η be the segment of $j_1\text{Ax}(w_1)$ where this fellow-traveling occurs. As $h_s\text{Ax}(v_s)$ may only 10δ -fellow-travel the first $105B$ units of γ_y , there must be a point r on η within $110B$ of x . Furthermore, r must lie within $110B$ of $\pi_{j_1\text{Ax}(w_1)}(x)$. We then have that a segment of either $[x, p_1]$ or $[x, p_2]$ must 10δ -fellow-travel η of length at least $150B$, contradicting our assumptions on $h\text{Ax}(v_i)$. As a consequence, we obtain the bounds

$$\text{diam}(\pi_{h_s\text{Ax}(v_s)}(j_1\text{Ax}(w_1))), \text{diam}(\pi_{j_1\text{Ax}(w_1)}(h_s\text{Ax}(v_s))) \leq 450B$$

Furthermore, from the fact that $h_s\text{Ax}(v_s)$ 10δ -fellow-travels γ_z and $j_1\text{Ax}(w_1)$ 10δ -fellow-travels γ_y , we must have that any geodesic from $\pi_{h_s\text{Ax}(v_s)}(j_1\text{Ax}(w_1))$ to $\pi_{j_1\text{Ax}(w_1)}(h_s\text{Ax}(v_s))$ passes within $900B$ of x . We apply Lemma 5.11 to conclude.

To obtain the first conclusion for $\bar{w}$, we note that we may find respective lifts z' , y' , and w' of $\bar{z}$, $\bar{y}$, and $\bar{w}$ so that w' and z' are minimal, z' and w' are near-minimal, and y'

and w' are near-minimal. From Lemma 6.1, $d_X(y', w') \leq 5000B$, while from the analysis above we have that any $[z', y']$ passes within $1000B$ of some lift x' of $\bar{x}$. From δ -thinness of triangles in X , $[w', z']$ must pass within $1200B$ of x' . The proof of the second condition is similar. $\square$

In order to make use of this lemma, we must demonstrate that we can find axes that are sufficiently ‘repulsing’ from a given set of enemy prongs. To do so, we use the notion of a *Schottky set*. These sets were first defined in [Bou+23], although the formulation we use is slightly modified from the formulation used by Gou  zel in [Gou22b]. Intuitively, given two independent loxodromic isometries in G , the construction gives a method for building a large number of distinct loxodromic isometries, all pointing in ‘different directions’ from a basepoint x_0 .

Definition 6.6. Let $\eta, E_G, E_D, E'_D, E_M \geq 0$. A finite set $\mathcal{E} \subseteq G$ of isometries is a $(\eta, E_G, E_D, E'_D, E_M)$ -*Schottky set* if the following conditions are satisfied:

1. For all $x, y \in X$, we have $|\{s \in \mathcal{E} : (x \cdot sy)_{x_0} \leq E_G\}| \geq (1 - \eta)|\mathcal{E}|$ and $|\{s \in \mathcal{E} : (x, s^{-1}y)_{x_0} \leq E_G\}| \geq (1 - \eta)|\mathcal{E}|$.
2. For all $s \in \mathcal{E}$, we have $E_D \leq d_X(x_0, sx_0) \leq E'_D$, and $d_G(e, s) \leq E_M$.

The following lemma is a straightforward consequence of the proof of Prop. 3.12 in [Gou22b]. It will be used to build axes in the main construction so that Lemma 6.5 can be applied.

Lemma 6.7 (Prop. 3.12, [Gou22b]). *For all $\eta > 0$ and $E, F > 0$, there exists E_G, E_D, E'_D and E_M and a set $\mathcal{E}$ such that $\mathcal{E}$ is a $(\eta, E_G, E_D, E'_D, E_M)$ -Schottky set, and such that $E_D \geq 10E_G + F$ and $E_G \geq 1000(\delta + E)$. Furthermore, E_G, E_D, E'_D and E_M depend only on η and the action $G \curvearrowright X$, and in particular are independent of n .*

7 Main construction

In this section, we aim to prove the following result. Once this result is proven, linear connectivity of the boundary will follow from standard arguments.

As usual, let n be a sufficiently large number, let $w_n^{(1)}, \dots, w_n^{(k)}$ be the n -step elements of independent permissible random walks on G and let $N = \langle\langle w_n^{(1)}, \dots, w_n^{(k)} \rangle\rangle$. We assume we are in the setup of Hypothesis 3.1, with M as in this hypothesis. Note that we may take M to be as large as needed by taking n sufficiently large.

We will fix a uniform basepoint $\bar{x}_0 \in X/N$ throughout these arguments. Before we begin, we note that by Lemma 4.5, all boundary points in $\partial(X/N)$ may be represented by $(1, 20L)$ -quasi-geodesic rays, and furthermore, when calculating Gromov products for boundary points as in Equation 1, it suffices to take the infimum over all representative sequences derived from $(1, 20L)$ -quasi-geodesic rays.




Figure 7: The construction used in Lemma 6.5, in the case where $\mathcal{S}(y, x)$ and $\mathcal{S}(x, z)$ each contain only one strut. The key point is that our assumptions force $\overline{j_1 \text{Ax}(w_1)}$ to ‘repulse’ from $[\bar{x}, \bar{p}_1]$ and $[\bar{x}, \bar{p}_2]$. As a result $\overline{j_1 \text{Ax}(w_1)}$ and $\overline{h_s \text{Ax}(v_s)}$ may only overlap on the scale of B (which does not grow with n), not on the scale of $\mathcal{M}$. This overlap bound plus properties of strut sequences forces the desired routing.

Our goal is to prove the following lemma, which will be broken up over several subsections and sublemmas.

Lemma 7.1. *Fix $\xi_1, \xi_2 \in \partial(X/N)$ with $(\xi_1 \cdot \xi_2)_{\bar{x}_0} = D$, along with a constant F . Then for all n sufficiently large in terms of F , the following property holds: there exists a finite sequence of boundary points $\xi_1 = \chi_1, \chi_2, \dots, \chi_Q = \xi_2$ such that for $1 \leq i \leq Q-1$, $(\chi_i \cdot \chi_{i+1})_{\bar{x}_0} \geq D + \frac{1}{2}F$, and for all $2 \leq i \leq Q-1$ and $j \in \{1, Q\}$, $(\chi_i \cdot \chi_j)_{\bar{x}_0} \geq D - 100L$.*

7.1 Building the Construction




Figure 8: A schematic of the construction used to prove Lemma 7.1, showing the relevant objects under projection to X/N . Note that the axes $\bar{A}_i$ are constructed using the points $\bar{u}_i$ as guides, and in particular must be repulsed from the points $\bar{y}_{i,j}$.

We will first set up the construction. Note that most relevant features of the construction will appear in Figure 8; keeping this figure in hand throughout the construction and its verification may be useful. We will also need several constants with a somewhat complex relationship; see Figure 9 as a guide.

Let $\bar{r}_1, \bar{r}_2 : \mathbb{Z}_{\geq 0} \rightarrow X/N$ be $(1, 20L)$ -quasi-geodesic rays originating at $\bar{x}_0$ and respectively representing ξ_1 and ξ_2 . Recalling that

$$L := \max_{hAx(v) \in \mathcal{A}} L(v),$$




Figure 9: A dependence chart of constants used in the main construction. An arrow from a to b indicates that b must depend on the value of a . An arrow from $G \curvearrowright X$ indicates constants that depend only on the action of G on X , including constants intrinsic to X . This graph must contain no cycles. We note that while F is an input to Lemma 7.1, we may need to replace it with a larger constant, and it is sufficient to verify the lemma for this new constant.

we let $\mathcal{Z}$ be a large constant, on the order of $1000L$ or larger, that will be chosen later. This constant is uniform with respect to ξ_1 , ξ_2 , and D . We may choose $\bar{r}_1$ and $\bar{r}_2$, and respective points $\bar{x}_1$ and $\bar{x}_2$ on these rays satisfying the following conditions:

1. We have $d_{X/N}(\bar{x}_0, \bar{x}_i) \geq \mathcal{Z}$.
2. We have $(\bar{x}_1, \bar{x}_2)_{\bar{x}_0} =_{10} D$.

Let x_0, x_1, x_2 be respective lifts of $\bar{x}_0, \bar{x}_1, \bar{x}_2$ to X , chosen via Lemma 5.10 so that x_1 and x_2 are minimal, x_0 and x_1 are near-minimal, and x_0 and x_2 are near-minimal. Let γ be a geodesic from x_1 to x_2 ; setting $\bar{\gamma} = \pi_{X/N}(\gamma)$, $\bar{\gamma}$ is a geodesic from $\bar{x}_1$ to $\bar{x}_2$. As X is action path-connected (see Definition 2.27), there exist a sequence of group elements $a_1, \dots, a_k$ such that $a_1 x_0 = x_1$, $a_k x_0 = x_2$, such that $d_G(a_i, a_{i+1}) \leq W$ for all $1 \leq i \leq k-1$, and so that $x_1 = a_1 x_0, a_2 x_0, \dots, a_{k-1} x_0, a_k x_0 = x_2$ is a unparameterized (E, E) -quasi-geodesic. Here, E and W are independent of n and L . We will fix the notation $w_i = a_i x_0$, and define $\bar{w}_i = \pi_{X/N}(w_i)$. Of course, $a_i \bar{x}_0 = \bar{w}_i$.

Now, there exists some constant E' depending only on E and δ so that every w_i lies within distance E' of γ , and hence every $\bar{w}_i$ lies within distance E' of $\bar{\gamma}$. Furthermore, for all $\bar{w}_i$ and $j \in \{1, 2\}$, we have $(\bar{w}_i \cdot \bar{x}_j)_{\bar{x}_0} \geq D - 10E'$. Note that this particular Gromov-product bound follows from the triangle inequality, and so holds in any geodesic metric space.

We now construct enemy prongs. Let B be the constant appearing in Lemma 6.5; we will need to assume this constant is sufficiently large in terms of several other constants, but not that it will *not* need to depend on n . For each w_i with $2 \leq i \leq k-1$, define $y_{i,1}, y_{i,2}$ via Lemma 6.2 so that, in the notation of that lemma, $y_1 = p_1(w_i, \bar{x}_0)$ and $y_2 = p_2(w_i, \bar{x}_0)$.

Lemma 7.2. *There exist constants E_G, E_D, E'_D and E_M satisfying*

$$E_D \geq 10(E_G + F + B), \quad E_G \geq 10000(\delta + F + E')$$

where E' is as in the preceding discussion, such that the following statement is true.

There exist group elements $s_i \in G$, points $u_i = w_i s_i x_0$ in X , and axes $A_i = h_i A x(v_i)$ (stretching notation so that here the v_i index is different from the index used in our hypotheses) such that for all $2 \leq i \leq k-1$ and $j \in \{1, 2\}$, the following properties hold, with appropriate modifications for $i = 2$ and $i = k-1$:

1. We have

$$(y_{i,j} \cdot u_i)_{w_i} \leq E_G, \quad (y_{i-1,j} \cdot u_i)_{w_i} \leq E_G, \quad \text{and} \quad (y_{i+1,j} \cdot u_i)_{w_i} \leq E_G.$$

2. We have

$$(u_{i-1} \cdot u_i)_{w_i} \leq E_G \quad \text{and} \quad (u_{i+1} \cdot u_i)_{w_i} \leq E_G.$$

3. We have

$$E_D \leq d_X(w_i, u_i) \leq E'_D.$$

4. We have

$$d_G(e, s_i) \leq E_M.$$

5. The points w_i and u_i are minimal. In particular,

$$d_{X/N}(\pi_N(w_i), \pi_N(u_i)) \geq E_D.$$

6. The axis A_i passes within distance 10δ of both u_i and u_{i+1} .

Proof. We apply Lemma 6.7. Consider w_2 . Set $\eta = \frac{1}{1000}$, and let E_G, E_D, E'_D, E_M be as in the aforementioned lemma, so that there exists a $(\eta, E_G, E_D, E'_D, E_M)$ -Schottky set $\mathcal{E}$ with respect to basepoint w_2 , and so that $E_D \geq 10E_G + F$ and $E_G \geq 1000(\delta + F + E')$, where E' is the constant from the action path-connectedness construction. Note that the dependence of these constants is somewhat tricky; see 9. However, note that these constants are independent of n and L . For $j \in \{1, 2\}$, by definition of $\mathcal{E}$,

$$|\{s \in \mathcal{E} : (y_{2,j} \cdot a_2 s x_0)_{w_2} \leq E_G\}| \geq \frac{999}{1000} |\mathcal{E}|$$

By a union bound, there exists at least one (in fact, many) $s_2 \in \mathcal{E}$ such that if $u_2 = a_2 s_2 x_0$, then $(y_{2,j} \cdot u_2)_{w_2} \leq E_G$ and $(y_{3,j} \cdot u_2)_{w_2} \leq E_G$ for $j \in \{1, 2\}$. Furthermore, we have $d_X(w_2, u_2) \geq E_D$, and $d_G(e, s_2) \leq E_M$.

For sufficiently large n , $E'_D \leq 0.01 \min_v(L(v))$, and we obtain as a result that for all $h\text{Ax}(v) \in \mathcal{A}$, $d_{h\text{Ax}(v)}(w_2, u_2) \leq 0.01L(v)$. So w_2 and u_2 have trivial strut sequence, and by Lemma 5.8 are minimal.

For each w_i up to w_{k-1} , we construct u_i and s_i analogously, except we also require that

$$(u_i \cdot u_{i-1})_{w_i} \leq E_G \text{ and } (u_i \cdot y_{i-1,j})_{w_i} \leq E_G.$$

These additional conditions pose no obstacle to the union bound argument. Notice the two Gromov products in the second item of the lemma statement are coarsely equal with coarseness only depending on the action $G \curvearrowright X$, so, up to allowing a slightly larger E_G , it suffices to prove only one of the two.

We now construct sequence of partially overlapping axes. Consider some u_i and u_{i+1} . Recalling $u_i = a_i s_i x_0$, we have

$$d_G(a_i s_i, a_{i+1} s_{i+1}) \leq d_G(e, s_i) + d_G(a_i, a_{i+1}) + d_G(e, s_{i+1}) \leq 2E_M + W$$

Notice that the bound $2E_M + W$ does not depend on n or L . We may therefore apply Theorem 2.26 and conclude that, for sufficiently large n , there exists an axis $A_i = h_i \text{Ax}(v_i)$ such that A_i passes within 10δ of both u_2 and u_3 (the word-length bound is not needed here). $\square$

Set $L_i = L(v_i)$. For each A_i , we choose two points $z_{i,1}, z_{i,2}$ on A_i such that $d_{A_i}(w_i, z_{i,1}) = 100\delta \cdot 0.4L(A_i)$, $d_{A_i}(w_i, z_{i,2}) = 100\delta \cdot 0.4L(A_i)$, and $d_{A_i}(z_{i,1}, z_{i,2}) = 100\delta \cdot 0.2L(A_i)$. Necessarily $d_X(u_i, z_{i,1}) \leq 0.45L(A_i)$, and $d_X(u_{i+1}, z_{i,2}) \leq 0.45L(A_i)$.

For each $z_{i,j}$, say $z_{i,j} = h_{i,j}x_0$.

Lemma 7.3. *For each point $z_{i,j}$ and axis $A_i = h_i Ax(v_i)$ as above, there exists a group element $c_{i,j} \in G$ with the following properties:*

1. *There exists a group element $c'_{i,j}$ so that $[x_0, c_{i,j}x_0] \subseteq \mathcal{N}_{20\delta}((c'_{i,j})^{-1}h_i Ax(v_i))$.*
2. *We have $d_X(x_0, c_{i,j}x_0) = 100\delta \cdot \frac{1}{3}L_i$.*
3. *No $(c'_{i,j})^{-1}h_i v_i h_i^{-1}c'_{i,j}$ -translate of $(c'_{i,j})^{-1}z_{i,j}$ lies within distance $\frac{1}{4}L_i$ of $[x_0, c_{i,j}x_0]$.*

Proof. Note that the $(c'_{i,j})^{-1}h_i v_i h_i^{-1}c'_{i,j}$ -translates of $(c'_{i,j})^{-1}z_{i,j}$ are spaced at distance at least $0.99L_i$ along $(c'_{i,j})^{-1}Ax(v_i)$. Therefore, we may construct the desired $c_{i,j}$ by choosing a segment of $(c'_{i,j})^{-1}Ax(v_i)$ with length approximately $\frac{1}{3}$ of a fully translation such that the segment is appropriately far away from any of these translates. $\square$

We now define a sequence $z_{i,j,q} = h_{i,j}c_{i,j}^q x_0$ for $q \in \mathbb{Z}_{\geq 0}$, and consider projections $\bar{z}_{i,j,q} = \pi_{X/N}(z_{i,j,q})$ of these points. We will define $\chi_{i,j}$ to be the boundary point derived from the sequence $(\bar{z}_{i,j,q})_q$. (Clearly we must check that this sequence defines a boundary point; we will do so shortly.) Our claim will be that the sequence of boundary points

$$\xi_1 = \chi_1, \chi_{2,1}, \chi_{2,2}, \chi_{3,1}, \dots, \chi_{k-1,1}, \chi_{k-1,2}, \chi_k = \xi_2$$

satisfies the conclusion of Lemma 7.1 after the appropriate reindexing.

7.2 Verifying the Construction

We have several claims to verify, which we split into a sequence of lemmas. As before, most relevant objects are depicted in Figure 8.

Lemma 7.4. *For all $2 \leq i \leq k-1$, $1 \leq j \leq 2$ the sequence $(\bar{z}_{i,j,q})_q$ defines a boundary point in $\partial(X/N)$, and in fact there exists a $(1, 1000M)$ -quasi-geodesic ray originating at $\bar{z}_{i,j,0}$ and passing through each $\bar{z}_{i,j,q}$.*

Proof. For ease of notation, fix i and j and set $z_q = z_{i,j,q}$ and $\bar{z}_q = \bar{z}_{i,j,q}$. Also, we may assume without loss of generality that $h_{i,j} = e$. We must show that

$$\lim_{q,r \rightarrow \infty} (\bar{z}_q \cdot \bar{z}_r)_{\bar{x}_0} = \infty$$

Consider the concatenation

$$[x_0, c_{i,j}x_0] \cup [c_{i,j}x_0, c_{i,j}^2x_0] \cup \dots \cup [c_{i,j}^{q-1}x_0, c_{i,j}^qx_0]$$

Each geodesic segment has length at least $0.31L_i$. On the other hand, observe that by construction

$$[c_{i,j}^n x_0, c_{i,j}^{n+1} x_0] \subseteq \mathcal{N}_{20\delta}(c_{i,j}^n (c'_{i,j})^{-1} h_i \text{Ax}(v_i)).$$

Therefore, if $[c_{i,j}^n x_0, c_{i,j}^{n+1} x_0]$ and $[c_{i,j}^{n+1} x_0, c_{i,j}^{n+2} x_0]$ 50δ -fellow-travel for more than distance $100\mathcal{M}$, then by the no-overlap condition, $c_{i,j}$ must fix the axis $(c'_{i,j})^{-1} h_i \text{Ax}(v_i)$. But then, as the length of $[x_0, c_{i,j} x_0]$ is approximately $\frac{1}{3}L_i$, we would find that the no-self-overlap condition is violated. In particular, $c_{i,j}$ does not fix $(c'_{i,j})^{-1} h_i \text{Ax}(v_i)$.

As $100\mathcal{M}$ is far smaller than $\frac{1}{3}L_i$, we conclude that the z_q are $200\mathcal{M}$ -aligned, and that any geodesic $[z_0, z_q]$ must decompose into segments that 50δ -fellow-travel translates of $\text{Ax}(v_i)$ for length at least $0.3L_i$ and at most $0.35L_i$, and ‘hops’ in between these segments having length at most $200\mathcal{M}$. It is straightforward to verify from these conditions that the resulting path is a $(1, 1000\mathcal{M})$ -quasi-geodesic.

We claim that all z_q are minimal with respect to z_0 . If not, there is some distinct z'_q that is minimal with respect to z_0 , with $z'_q = h z_q$ for $h \in N \setminus \{e\}$ and $d_X(z'_q, z_0) < d_X(z_q, z_0)$. By Lemma 4.1, there is some axis $h \text{Ax}(v)$ such that $d_{h \text{Ax}(v)}(z'_q, z_q) \geq 0.99L(v)$. As z_0 and z'_q are minimal, $d_{h \text{Ax}(v)}(z'_q, z_0) \leq 0.51L(v)$, so we must have $d_{h \text{Ax}(v)}(z_q, z_0) \geq 0.4L(v)$ by the triangle inequality. Therefore, there is a segment γ of $[z_0, z_q]$ that 50δ -fellow-travels $h \text{Ax}(v)$ for length at least $0.36L(v)$, and this segment must overlap a segment on which $[z_0, z_q]$ 50δ -fellow-travels a translate of $\text{Ax}(v_i)$, with the overlap having length at least $1000\mathcal{M}$. We have that $h \text{Ax}(v)$ cannot be equal to this translate, as the axis-fellow-traveling segments of $[z_0, z_q]$ have length uniformly bounded above by $0.35L_i$. Therefore, we obtain a match between $h \text{Ax}(v)$ and some translate of $\text{Ax}(v_i)$, which is a contradiction by the no-overlap condition.

By translation, we obtain that all elements of the sequence $(z_q)_q$ are pairwise minimal. Therefore $(\bar{z}_q \cdot \bar{z}_r)_{\bar{z}_0} = (z_q \cdot z_r)_{z_0}$ for all q, r . Furthermore, by the quasi-geodesic claim made above, it is more or less immediate that $\lim_{q,r \rightarrow \infty} (z_q \cdot z_r)_{z_0} = \infty$. $\square$

Lemma 7.5. *For all $z_{i,j,q}$ and all $h \text{Ax}(v) \in \mathcal{A}$, we have*

$$d_{h \text{Ax}(v)}(z_{i,j,q}, w_i) \leq 0.41L(v).$$

In particular, $z_{i,j,q}$ and w_i are minimal.

Furthermore, there exists a constant B independent of n and F with the following property. For M sufficiently large in terms of B , and for all $\bar{z}_{i,j,q}$ and any geodesic $[\bar{x}_0, \bar{z}_{i,j,q}]$, this geodesic passes within $1000B$ of $\bar{w}_i$ and within $150\mathcal{M}$ of $\bar{z}_{i,j,0}$.

Proof. It essentially suffices to verify the hypotheses of Lemma 6.5. Let w_i and $z_{i,j,q}$ be as in the construction, and recall that $y_{i,1}$ and $y_{i,2}$ are constructed as enemy prongs relative to $\bar{x}_0$. Assume that $j = 1$; the case for $j = 2$ is analogous.

First, we check that w_i and $z_{i,1,q}$ are minimal. Consider the path $[w_i, z_{i,1,0}] \cup [z_{i,1,0}, z_{i,1,q}]$. The axis A_i passes within 10δ of u_i and u_{i+1} , and by construction $(u_i, u_{i+1})_{w_i} \leq E_G$. Furthermore, both u_i and u_{i+1} have distance at least $9E_G$ from w_i , and so we may conclude

that A_i must pass within $2E_G$ of w_i . Furthermore, $z_{i,1,0}$ lies on A_i , and $d_{A_i}(w_i, z_{i,1,0}) = 100\delta \cdot 0.4L(A_i)$. We conclude by the Bounded Geodesic Image Theorem that $[w_i, z_{i,1,0}]$ has length at least $0.39L(A_i)$ and at most $0.41L(A_i)$, and 50δ -fellow-travels A_i except for a segment of length at most $4E_G$.

We verified in the proof of Lemma 7.4 that there is a $(1, 1000\mathcal{M})$ -quasi-geodesic path from $z_{i,1,0}$ to $z_{i,1,q}$, which 50δ -fellow-travels a sequence of distinct translates of A_i for lengths between $0.3L_i$ and $0.35L_i$, with segments in between of length at most $200\mathcal{M}$. We also know that $[z_{i,1,0}, z_{i,1,1}]$ 50δ -fellow-travels some translate A'_i of $\text{Ax}(v_i)$ for length at least $0.32L_i$ and at most $0.34L_i$, and that $[w_i, z_{i,1,0}]$ $5E_G$ -fellow-travels A_i for length at least $0.39L_i$ and at most $0.41L_i$.

We claim that $A'_i \neq A_i$, using Lemma 7.3. If not, then as

$$[h_{i,1}x_0, h_{i,1}c_{i,1}x_0] = [z_{i,1,0}, z_{i,1,1}] \subseteq \mathcal{N}_{20\delta}(h_{i,1}(c'_{i,1})^{-1}h_i\text{Ax}(v_i))$$

we must have that $A'_i = h_{i,1}(c'_{i,1})^{-1}h_i\text{Ax}(v_i)$, and as $h_{i,1}(c'_{i,1})^{-1}$ therefore fixes $A'_i = A_i = h_i\text{Ax}(v_i)$, we must have that $h_{i,j}(c'_{i,1})^{-1} \in \langle h_i v_i h_i^{-1} \rangle$. As a consequence of Lemma 7.3, no $h_{i,1}(c'_{i,1})^{-1}h_i v_i h_i^{-1} c'_{i,1}(h_{i,1})^{-1}$ -translate of $h_{i,1}(c'_{i,1})^{-1}z_{i,1,0}$ lies within distance $\frac{1}{4}L_i$ of $[z_{i,1,0}, z_{i,1,1}]$. Using $h_{i,j}(c'_{i,1})^{-1} \in \langle h_i v_i h_i^{-1} \rangle$ and the fact that $z_{i,1,0}$ lies on A_i , we have that no $h_i v_i h_i^{-1}$ -translate of $z_{i,1,0}$ lies within distance $\frac{1}{4}L_i$ of $[z_{i,1,0}, z_{i,1,1}]$, which is clearly contradictory.

As such, $[w_i, z_{i,1,0}]$ and $[z_{i,1,0}, z_{i,1,1}]$ cannot 50δ -fellow-travel for length more than $100\mathcal{M}$ without inducing a failure of the no-overlap condition. Arguing as in Lemma 7.4, we may patch axis-following segments together and conclude that for any $h\text{Ax}(v) \in \mathcal{A}$,

$$d_{h\text{Ax}(v)}(w_i, z_{i,1,q}) \leq 0.41L(v).$$

Therefore w_i is minimal relative to all $z_{i,1,q}$. As a further consequence, any geodesic $[w_i, z_{i,1,q}]$ must 50δ -fellow-travel $[w_i, z_{i,1,0}]$ until it has passed within $1000\mathcal{M}$ of $z_{i,1,0}$. Since w_i lies within $4E_G$ of $h_i\text{Ax}(v_i)$ and $z_{i,1,0}$ lies on $h_i\text{Ax}(v_i)$, we must have that $h_i\text{Ax}(v_i)$ is the first strut appearing in $\mathcal{S}(w_i, z_{i,1,q})$. By analogous analysis, the axis that 50δ -fellow-travels $[z_{i,1,0}, z_{i,1,1}]$ is the second strut appearing in $\mathcal{S}(w_i, z_{i,1,q})$, and if this axis is labeled $h'_i\text{Ax}(v'_i)$, $z_{i,1,0} \in \mathcal{N}_{120\mathcal{M}}(\pi_{h_i\text{Ax}(v_i)}(h'_i\text{Ax}(v'_i)))$. We aim to apply this analysis together with Lemma 6.5 to conclude that for any near-minimal respective lifts x'_0 of x_0 and $z'_{i,1,q}$ of $\bar{z}_{i,1,q}$, any $[x'_0, z'_{i,1,q}]$ passes within $150\mathcal{M}$ of some lift of $\bar{z}_{i,1,0}$.

Next, it follows from the geometric analysis performed above, the fact that $d_X(w_i, u_i) \geq 9E_G$, and a standard thin-quadrilaterals argument that any geodesic $[w_i, z_{i,1,q}]$ passes within 50δ of u_i . Recall that for $j \in \{1, 2\}$, $(y_{i,j} \cdot u_i)_{w_i} \leq E_G$. Therefore, for B sufficiently large in terms of E_G , we have that $(y_{i,1} \cdot z_{i,j,q})_{w_i}$ and $(y_{i,2} \cdot z_{i,j,q})_{w_i}$ are both bounded above by B . Lastly, assume we have some axis $h\text{Ax}(v)$ such that $h\text{Ax}(v)$ 50δ -fellow-travels $[w_i, z_{i,j,q}]$ for length at least $0.0001L(v)$, and that this axis $10B$ -fellow-travels, say, $[w_i, y_{i,1}]$ for distance at least $200B$. By taking M and L sufficiently large, we may assume that

$\mathcal{M} \geq 10^6 * B$, and so $hAx(v)$ must pass within $10\mathcal{M}$ of w_i . From the structure of the $(1, 1000\mathcal{M})$ -quasi-geodesic constructed above between w_i and $z_{i,j,q}$, we have that $hAx(v)$ must 50δ -fellow-travel one of the axes tracked by this quasi-geodesic for length at least $0.00001L(v)$, and thus be equal to one of these axes (as otherwise the no-overlap condition would be violated). But then its proximity to w_i forces $hAx(v) = A_i$. Lastly, A_i passes within 10δ of both u_i and u_{i+1} , while $(y_{i,1} \cdot u_i)_{w_i}$ and $(y_{i,1} \cdot u_{i+1})_{w_{i+1}}$ are bounded above by E_G . Recall that $d_X(w_i, w_{i+1}) \leq 10E$ (where E is the action path-connectedness constant), so for E_G taken suitably large relative to E , $(y_{i,1} \cdot u_{i+1})_{w_i} \leq 2E_G$ as well. For B sufficiently large in terms of E_G , a straightforward δ -hyperbolicity argument now shows that A_i cannot $10B$ -fellow-travel $[w_i, y_{i,1}]$ for distance more than $200B$.

Lemma 6.5 now yields both results. $\square$

We now turn to the Gromov-product claims.

Lemma 7.6. *For all $1 \leq i \leq Q-1$, we have $(\xi_1 \cdot \chi_i)_{\bar{x}_0} \geq D-100L$ and $(\xi_2, \chi_i)_{\bar{x}_0} \geq D-100L$.*

Proof. We prove the case of ξ_1 ; the case of ξ_2 follows by identical argument. Let $\bar{x}'_1$ be a point on $\bar{r}_1$; we may assume by taking $\bar{x}'_1$ sufficiently far from $\bar{x}_0$ that $d_{X/N}(\bar{x}_0, \bar{x}'_1) \geq 2d_{X/N}(\bar{x}_0, \bar{x}_1)$. Applying the triangle inequality for Gromov products together with Lemma 4.5, for any $z_{i,j,q}$,

$$(\bar{x}'_1 \cdot \bar{z}_{i,j,q})_{\bar{x}_0} \geq \min((\bar{x}'_1, \bar{w}_i)_{\bar{x}_0}, (\bar{w}_i, \bar{z}_{i,j,q})_{\bar{x}_0}) - 10L$$

By Lemma 7.5 and the previous observation on $d_{X/N}(\bar{w}_i, \bar{x}_0)$,

$$(\bar{w}_i, \bar{z}_{i,j,q})_{\bar{x}_0} \geq d_{X/N}(\bar{w}_i, \bar{x}_0) - 2000B \geq D - 10E' - 3000B$$

On the other hand,

$$(\bar{x}'_1, \bar{w}_i)_{\bar{x}_0} \geq \min((\bar{x}'_1 \cdot \bar{x}_1)_{\bar{x}_0}, (\bar{x}_1 \cdot \bar{w}_i)_{\bar{x}_0}) - 10L$$

By setting $\mathcal{Z}$ (which controls $d_{X/N}(\bar{x}_0, \bar{x}_1)$) sufficiently large in terms of L and D , we can ensure $(\bar{x}'_1 \cdot \bar{x}_1)_{\bar{x}_0} \geq 10D$, as $\bar{x}'_1$ and $\bar{x}_1$ lie on the same quasi-geodesic ray $\bar{r}_1$. On the other hand, $(\bar{x}_1 \cdot \bar{w}_i)_{\bar{x}_0}$ is at least $D - 1000B$ by Lemma 7.5 and the triangle inequality. After unwinding the minima, we obtain that for the particular representatives of ξ_1 and χ_i chosen, the limiting Gromov product is at least $D - 25L$. We may now apply Lemma 2.18 to conclude. $\square$

Lemma 7.7. *For all χ_i , $1 \leq i \leq Q-1$, $(\chi_i \cdot \chi_{i+1})_{\bar{x}_0} \geq D + \frac{1}{2}F$.*

Proof. Without loss of generality, we assume that $F \geq 10^9B$, as B does not depend on F . We then require that $\mathcal{Z} \geq D + F + 10^6 * L$. Recalling the reindexing required to construct $\chi_1, \dots, \chi_Q$, there are a few cases to handle.

The first case is ξ_1 and $\chi_{2,1}$. For this case, Lemma 7.5 implies that any geodesic from $\bar{z}_{2,1,q}$ to $\bar{x}_0$ passes within $1000B$ of $\bar{w}_2$, and by construction $(\bar{x}_1, \bar{w}_2)_{\bar{x}_0} \geq \mathcal{Z} - 10L$. Arguing as in Lemma 7.6 then yields that $(\xi_1, \chi_{2,1})_{\bar{x}_0} \geq \mathcal{Z} - 100L$.

The case of ξ_2 and $\chi_{k-1,2}$ is analogous.

Next, we consider $\chi_{i,1}$ and $\chi_{i,2}$. Let $\bar{\rho}_1$ be a $(1, 20L)$ -quasi-geodesic ray representing $\chi_{i,1}$, and let $\bar{\rho}_2$ be a $(1, 20L)$ -quasi-geodesic ray representing $\chi_{i,2}$, such that both rays originate from $\bar{x}_0$. Recall the existing $(1, 1000\mathcal{M})$ -quasi-geodesic rays representing $\chi_{i,1}$ and $\chi_{i,2}$ constructed from the $z_{i,j,q}$ points.

Applying the Morse Lemma, every point $\bar{s}_1$ sufficiently far along on $\bar{\rho}_1$ satisfies $d_{X/N}(\bar{x}_0, \bar{s}_1) \geq 10^6 * L$, $d_{X/N}(\bar{w}_i, \bar{s}_1) \geq 10^6 * L$, $d_{X/N}(\bar{z}_{i,1,0}, \bar{s}_1) \geq 10^6 * L$, and $d_{X/N}(\bar{s}_1, \bar{z}_{i,1,q}) \leq 200L$ for some $\bar{z}_{i,1,q}$. Applying Lemma 5.10, there exists s_1 lifting $\bar{s}_1$ and x'_0 lifting $\bar{x}_0$ such that s_1 and x_0 are minimal, $z_{i,1,q}$ and x_0 are near-minimal, and $z_{i,1,q}$ and s_1 are near-minimal. Lemma 6.1 then implies $d_X(s_1, z_{i,1,q}) \leq 2000L$.

By Lemma 7.5, $[z_{i,1,q}, x'_0]$ passes within $1000B$ of some lift w'_i of $\bar{w}_i$, and within $150\mathcal{M}$ of some lift $z'_{i,1,0}$ of $\bar{z}_{i,1,0}$. As $z_{i,1,q}$ and s_1 both have distance at least 10^5L from w'_i and $z'_{i,1,0}$, we may conclude by δ -thinness of the triangle with sides $[s_1, x'_0]$, $[z'_{i,1,q}, x'_0]$, and $[s_1, z'_{i,1,q}]$, along with minimality of s_1 and x'_0 , that $[\bar{x}_0, \bar{s}_1]$ passes within $1100B$ of $\bar{w}_i$ at some point $\bar{s}_1^w$ and within $160\mathcal{M}$ of $\bar{z}_{i,1,0}$ at some point $\bar{s}_1^z$. Let s_1^w be the unique lift of $\bar{s}_1^w$ that lies within $1100B$ of w_i , and let s_1^z be the unique lift of $\bar{s}_1^z$ that lies within $160\mathcal{M}$ of $z_{i,1,0}$. Essentially by the same argument used to prove that w_i and $z_{i,1,0}$ are minimal in Lemma 7.5, s_1^w and s_1^z are minimal. By a standard thin-quadrilaterals argument, any geodesic $[s_1^w, s_1^z]$ passes through some point t_1 lying within distance 50δ of A_i and within distance $40000\mathcal{M}$ of $z_{i,1,0}$. Let $\bar{t}_1$ be the projection of such a point on $[\bar{x}_0, \bar{s}_1]$.

By identical argument, $[\bar{x}_0, \bar{s}_2]$ passes through some point $\bar{t}_2$ lying within distance 50δ of $\bar{A}_i$ and within distance $40000\mathcal{M}$ of $\bar{z}_{i,1,0}$.

From here, we observe the following inequality, which follows solely from the triangle inequality and does not require hyperbolicity:

$$(\bar{s}_1 \cdot \bar{s}_2)_{\bar{x}_0} \geq (\bar{t}_1 \cdot \bar{t}_2)_{\bar{x}_0}.$$

Using the lifts t_1 of $\bar{t}_1$ and t_2 of $\bar{t}_2$ constructed above, we have that $d_{A_i}(t_j, w_i)$ both lie in $[0.39L(A_i), 0.41L(A_i)]$ for $j \in \{1, 2\}$. As a consequence, $d_{X/N}(\bar{t}_1, \bar{t}_2) \leq 0.25L(A_i)$. Now, it is straightforward to verify as in Lemma 7.5 that t_1 and t_2 satisfy the hypotheses of Lemma 6.5 (by the same argument used for $z_{i,1,0}$), and so $[\bar{t}_1, \bar{x}_0]$ and $[\bar{t}_2, \bar{x}_0]$ both pass within $1000B$ of $\bar{w}_i$, say at points $\bar{t}'_1$ and $\bar{t}'_2$. By taking the (unique) lifts of these points that lie within distance $1000B$ close to w_i , we may check minimality as above and so verify that $d_{X/N}(\bar{t}_1, \bar{t}'_1)$ and $d_{X/N}(\bar{t}_2, \bar{t}'_2)$ are both at least $0.38L(A_i)$. Therefore, for $j \in \{1, 2\}$,

$$d_{X/N}(\bar{t}_j, \bar{x}_0) \geq 0.37L(A_i) + D$$

and so

$$(\bar{s}_1, \bar{s}_2)_{\bar{x}_0} \geq (\bar{t}_1, \bar{t}_2)_{\bar{x}_0} = \frac{d_{X/N}(\bar{t}_1, \bar{x}_0) + d_{X/N}(\bar{t}_2, \bar{x}_0) - d_{X/N}(\bar{t}_1, \bar{t}_2)}{2}.$$

$$\geq \frac{0.74L(A_i) + 2D - 0.25L(A_i)}{2} \geq D + F$$

as long as M (and hence $\min_i L_i$) is chosen to be sufficiently large with respect to F . As $\bar{s}_1$ and $\bar{s}_2$ were chosen arbitrarily, we obtain the conclusion that $(\chi_{i,1} \cdot \chi_{i,2})_{\bar{x}_0} \geq D + F$ (and in fact obtain a significantly larger increment).

The last remaining case is $\chi_{i,2}$ and $\chi_{i+1,1}$. Copying the argument used in the previous case, we obtain $\bar{t}_1$ and $\bar{t}_2$ respectively lying close to $\bar{z}_{i,1,0}$ and $\bar{z}_{i,2,0}$. Furthermore, $[\bar{t}_1, \bar{x}_0]$ must pass within $1000B$ of $\bar{w}_{i+1}$, say at some point $\bar{t}'_1$, and $[\bar{t}_2, \bar{x}_0]$ must pass within $1000B$ of $\bar{w}_{i+1}$, say at some point $\bar{t}'_2$. Let t'_1 and t'_2 be respective lifts of these points lying within $1000B$ of $\bar{w}_{i+1}$. Now, t_1 , t'_1 , w_{i+1} , and u_{i+1} all lie within $200M$ of a subsegment of A_i with length at most $0.41L_i$, and it is straightforward to check from this fact that these four points are pairwise minimal. Recall that u_{i+1} lies within 10δ of A_i , and its projection to A_i lies between the projection of w_i and the projection of $z_{i,2,0}$, with an analogous statement for u_{i+1} and A_{i+1} . As $d_{X/N}(\bar{w}_{i+1}, \bar{u}_{i+1}) \geq 10E_G + F$ and F be chosen to be sufficiently large with respect to B , we may assert that any $[\bar{t}_1, \bar{x}_0]$ passes within 100δ of $\bar{u}_{i+1}$, say at some point $\bar{y}_1$. Likewise, $[\bar{t}_2, \bar{x}_0]$ passes within 100δ of $\bar{u}_{i+1}$, say at $\bar{y}_2$ - here it is important that both A_i and A_{i+1} were constructed with u_{i+1} as a 'guide'. Lastly, we may check as in Lemma 7.5 that any geodesic $[\bar{u}_{i+1}, \bar{x}_0]$ or $[\bar{y}_1, \bar{x}_0]$ passes within $1000B$ of $\bar{w}_i$. As a result,

$$d_{X/N}(\bar{u}_{i+1}, \bar{x}_0) \geq d_{X/N}(\bar{u}_{i+1}, \bar{w}_{i+1}) + d_{X/N}(\bar{w}_{i+1}, \bar{x}_0) - 10000B$$

We already know that $d_{X/N}(\bar{w}_{i+1}, \bar{x}_0) \geq D - B$ for B sufficiently large in terms of E and E' , as well as $d_{X/N}(\bar{w}_{i+1}, \bar{u}_{i+1}) \geq 10E_G + F$ by the minimality claim of Lemma 7.2. We conclude that

$$\begin{aligned} (\bar{s}_1 \cdot \bar{s}_2)_{\bar{x}_0} &\geq (\bar{t}_1 \cdot \bar{t}_2)_{\bar{x}_0} \geq (\bar{u}_{i+1} \cdot \bar{u}_{i+1})_{\bar{x}_0} - 1000\delta \\ &= d_{X/N}(\bar{u}_{i+1}, \bar{x}_0) - 1000\delta \geq F + D - 20000B \end{aligned}$$

Choosing F sufficiently large in terms of B produces the desired bound $(\chi_{i,2} \cdot \chi_{i+1,1})_{\bar{x}_0} \geq D + \frac{1}{2}F$. $\square$

We have completed the proof of Lemma 7.1.

7.3 Linear Connectivity

We obtain linear connectivity of the boundary as a result of Lemma 7.1.

Theorem 7.8. *Assume that $G \curvearrowright X$ and random elements $w_1, \dots, w_k$ chosen as the n -step element of independent permissible random walks satisfy Hypothesis 3.1, where n is sufficiently large in terms of $G \curvearrowright X$. Then there exists a constant C' such that for all $\xi_0, \xi_1 \in \partial(X/N)$, there exists a path $\gamma(t) : [0, 1] \rightarrow \partial(X/N)$ with $\gamma(0) = \xi_0$, $\gamma(1) = \xi_1$, and for all $0 \leq t \leq 1$ and $j \in \{0, 1\}$,*

$$(\xi_j \cdot \gamma(t))_{\bar{x}_0} \geq (\xi_0 \cdot \xi_1)_{\bar{x}_0} - C'.$$

As an immediate consequence, $\partial(X/N)$ is linearly connected.

Proof. Given Lemma 7.1, the remainder of the proof is essentially standard, and we will follow the proof of Prop. 2.2 in [Wri24]. We present the details for clarity.

Recall that the visual metric d_v on $\partial(X/N)$ satisfies the property that there exists $a > 1$ and $C > 1$ so that for all $\xi, \chi \in \partial(X/N)$,

$$\frac{1}{C}a^{-(\xi \cdot \chi)_{\bar{x}_0}} \leq d_v(\xi, \chi) \leq Ca^{-(\xi \cdot \chi)_{\bar{x}_0}}.$$

Given points ξ_0 and ξ_1 on $\partial(X/N)$, we will iteratively define functions $\gamma_r : I_r \rightarrow \partial(X/N)$, for $r \in \mathbb{Z}_{\geq 0}$, where $I_r \subseteq [0, 1]$ and $I_r \subseteq I_{r+1}$ for all r . We further require that for all $r < r'$, $\gamma_{r'}|_{I_r} = \gamma_r$. We will set the convention that $\xi_t := \gamma_r(t)$ for any $t \in [0, 1]$ and any r such that $t \in I_r$. The base case is $I_0 = \{0, 1\}$ and $\gamma_0(0) = \xi_0$, $\gamma_0(1) = \xi_1$.

Define $(\xi_0 \cdot \xi_1)_{\bar{x}_0} = A$.

We make the following inductive hypotheses:

1. We have $I_r \subseteq I_{r+1}$, and $\gamma_{r+1}|_{I_r} = \gamma_r$.
2. If t_1 and t_2 are adjacent points of I_r with $t_1 < t_2$, then

$$(\xi_{t_1} \cdot \xi_{t_2})_{\bar{x}_0} \geq A + \frac{1}{2}Fr$$

3. If t_1 and t_3 are adjacent points of I_r with $t_1 < t_3$, and $t_2 \in I_{r+1}$ satisfies $t_1 \leq t_2 \leq t_3$, then

$$(\xi_{t_1} \cdot \xi_{t_2})_{\bar{x}_0} \geq A + \frac{1}{2}Fr - 100L$$

Assuming these hypotheses are satisfied, the inductive procedure consists of applying Lemma 7.1 to each pair of adjacent points of I_r . Specifically, say t_1 and t_2 are adjacent in I_r with $t_1 < t_2$, so $(\xi_{t_1} \cdot \xi_{t_2})_{\bar{x}_0} \geq A + \frac{1}{2}Fr$. Lemma 7.1 gives a list $\xi_{t_1} = \chi_0, \chi_1, \dots, \chi_Q = \xi_{t_2}$ such that for $1 \leq i \leq Q - 1$, $(\chi_i \cdot \chi_{i+1})_{\bar{x}_0} \geq A + \frac{1}{2}F(r + 1)$, while for all $1 \leq i \leq Q - 1$, $(\chi_0 \cdot \chi_i)_{\bar{x}_0} \geq A + \frac{1}{2}Fr - 100L$. We may therefore set

$$I_{r+1} \cap [t_1, t_2] = \left\{ t_1 + i \frac{t_2 - t_1}{n} : i \in \mathbb{Z} \cap [0, Q] \right\}$$

and define $\gamma_{r+1}\left(t_1 + i \frac{t_2 - t_1}{n}\right) = \chi_i$. Performing this construction for all pairs of consecutive points in I_r yields I_{r+1} satisfying the inductive hypotheses.

We set $I = \bigcup_{r=0}^{\infty} I_r$, and may by the consistency assumption define $\gamma : I \rightarrow \partial X$ by extending the γ_r . Next, we observe the formula

$$\text{diam}(\gamma(I)) \leq \text{diam}(\gamma(I_0)) + 2 \sum_{r=0}^{\infty} \max_{s \in I_{r+1}} \min_{t \in I_r} d_v(\gamma(s), \gamma(t))$$

From our inductive hypotheses, for every $s \in I_{r+1}$, there is a $t \in I_r$ with $(\xi_s \cdot \xi_t)_{\bar{x}_0} \geq A + Fr - 100L$. Translating to the visual metric,

$$d_v(\gamma(s), \gamma(t)) \leq Ca^{-A-Fr+100L}$$

so we obtain

$$\text{diam}(\gamma(I)) \leq Ca^{-A} + 2C \sum_{r=0}^{\infty} a^{-A-Fr+100L}.$$

This expression is certainly bounded by $a^{-A} \geq \frac{1}{C} d_v(\xi_0, \xi_1)$ times a constant $\mathcal{L}$ depending only on C and L . We conclude that there exists some constant C' such that for that every $t \in I$,

$$(\xi_1 \cdot \gamma(t))_{\bar{x}_0} \geq (\xi_1 \cdot \xi_2)_{\bar{x}_0} - C'$$

Essentially the same estimate yields that if t_1 and t_2 are consecutive points of I_r , then the diameter of $\gamma(I \cap [t_1, t_2])$ is at most $a^{-A}\mathcal{L}$ times a^{-Fr} . Given any $\varepsilon > 0$, we may therefore find R such that if t_1 and t_2 are consecutive points on I_R , then $\text{diam}(\gamma(I \cap [t_1, t_2])) \leq \varepsilon/2$. Choosing δ to be, say, $\frac{1}{10}$ of the length of the smallest distance between consecutive points of I_R , then for any s_1, s_2 on I with $|s_1 - s_2| < \delta$, there is a $z \in I_R$ within δ of both x and y . From the triangle inequality between $\gamma(s_1)$, $\gamma(s_2)$, and $\gamma(z)$, we have $d(\gamma(s_1), \gamma(s_2)) \leq \varepsilon$.

Therefore, γ is uniformly continuous on I , and must extend to a continuous function on the closure $\bar{I}$. From the iterative construction, it is straightforward to check that $\bar{I}$ must equal $[0, 1]$, and so $\gamma([0, 1])$ is connected. Lastly, we have that the diameter of $\gamma(\bar{I}) = \gamma(I)$ is equal to the diameter of $\gamma(I)$.

We conclude that $\partial(X/N)$ is linearly connected. $\square$

Proof of Theorem 1.1. With the setup of Theorem 1.1, we have that Theorem 2.26 implies that Hypothesis 3.1 holds with probability approaching 1 as $n \rightarrow \infty$. Theorem 7.8 now supplies the desired result. $\square$

8 Dimension of the Boundary

Remark 8.1. In this section, all topological spaces will be assumed to be separable metric spaces. This condition holds for all spaces relevant to us, and avoids some edge cases related to the definition of topological dimension.

8.1 Goal and Setup

Recall the following definition.

Definition 8.2. A separable metric space S has topological dimension $\dim(S) = k$ if k is the minimum nonnegative integer such that every open cover $\mathcal{U}$ of S admits a refinement $\mathcal{U}'$ such that every $s \in S$ lies in at most $k + 1$ elements of $\mathcal{U}'$.

We note that for all $T \subseteq S$ with the subspace topology, $\dim(T) \leq \dim(S)$. Furthermore, if S is path-connected, then $\dim(S) \geq 1$.

Our goal is to prove the following lemma.

Lemma 8.3. *Let G , X , and N be as in Hypothesis 3.1. Then*

$$1 \leq \dim(\partial(X/N)) \leq \dim(\partial X) + 1.$$

Before going on, we state two topological results. Both results are classical.

Lemma 8.4. *Let X be a separable metric space, and let $T_1, T_2, \dots$ be a countable collection of closed subsets. Then*

$$\dim\left(\bigcup_{i=1}^{\infty} T_i\right) \leq \max_i \dim(T_i)$$

where all subsets carry the subspace topology.

Lemma 8.5 (Urysohn-Menger Sum Formula). *Let X be a separable metric space, and let Y_1 and Y_2 be arbitrary subspaces of X such that $X = Y_1 \cup Y_2$. Then*

$$\dim(X) \leq \dim(Y_1) + \dim(Y_2) + 1$$

where Y_1 and Y_2 carry the subspace topology.

We now cover a few elementary results concerning boundary sets of hyperbolic spaces.

Definition 8.6. Let X be a geodesic metric space and let A and B be two subsets of X with $A \subseteq B$. We define the *star* of B relative to A , $B^*(A)$, as

$$B^*(A) = \{x \in X : \exists a \in A, b \in B \text{ such that } x \in [a, b]\}.$$

Definition 8.7. Let X be a Gromov hyperbolic metric space, and let A be a subset of M . We define the *limit set* $\Lambda(A)$ as the subset of ∂M definable via sequences lying entirely in A .

The following two lemmas are elementary.

Lemma 8.8. *Let X be a Gromov hyperbolic geodesic metric space, and let $A \subseteq B \subseteq X$ with $A \neq \emptyset$ and $\text{diam}(A) < \infty$. Then $B^*(A)$ is quasi-convex.*

Lemma 8.9. *Let X be a Gromov hyperbolic simplicial graph, and let A be a connected quasi-convex subgraph of X . Equip A with the induced subgraph metric. Then $\Lambda(A)$ is naturally identified with ∂A . Furthermore, $\Lambda(A)$ is closed in the visual-metric topology on ∂X , and the intrinsic visual-metric topology on ∂A agrees with the subspace topology on $\Lambda(A)$.*

8.2 Strut Prefixes and Decomposition

Throughout, we fix an arbitrary basepoint $\bar{x}_0 \in X$. Recall that $\mathcal{A}$ is the set of all axes $h\text{Ax}(v)$ in X , and $\mathcal{A}/N$ is the set of all images of axes $\overline{h\text{Ax}(v)}$ in X/N .

Definition 8.10. Given $\bar{x}, \bar{y} \in X/N$, choose near-minimal respective lifts x and y of $\bar{x}$ and $\bar{y}$, and enumerate $\mathcal{S}(x, y)$ as

$$x <_{\mathcal{S}} h_1\text{Ax}(v_1) <_{\mathcal{S}} \dots <_{\mathcal{S}} h_k\text{Ax}(v_k) <_{\mathcal{S}} y$$

We then define the *quotient strut sequence* $\overline{\mathcal{S}}(\bar{x}, \bar{y})$ to be the ordered sequence

$$\bar{x} <_{\overline{\mathcal{S}}} \overline{h_1\text{Ax}(v_1)} <_{\overline{\mathcal{S}}} \dots <_{\overline{\mathcal{S}}} \overline{h_k\text{Ax}(v_k)} <_{\overline{\mathcal{S}}} \bar{y}.$$

By Lemma 5.11, $\overline{\mathcal{S}}(\bar{x}, \bar{y})$ is well-defined, regardless of the lifts chosen.

Definition 8.11. We define the *non-turning locus based at $\bar{z} \in S'_{\bar{z}}$* as the set of all points $\bar{x} \in X/N$ such that there exist minimal respective lifts x and z of $\bar{x}$ and $\bar{z}$ such that for all axes $h\text{Ax}(v) \in \mathcal{A}$,

$$D_{h\text{Ax}(v)}(x, z) \leq 0.39L(v).$$

We then define the set $S_{\bar{z}}$ by

$$S_{\bar{z}} = (S'_{\bar{z}})^*(\{z\}).$$

Lemma 8.12. *Let $\bar{x} \in S_{\bar{x}_0}$. Then there exists minimal respective lifts x and x_0 of $\bar{x}$ and $\bar{x}_0$ such that for all axes $h\text{Ax}(v) \in \mathcal{A}$,*

$$D_{h\text{Ax}(v)}(x, x_0) \leq 0.4L(v).$$

Proof. For any $\bar{y} \in S_{\bar{x}_0}$, there exists some $\bar{x} \in X/N$ and a geodesic $[\bar{x}_0, \bar{x}]$ such that $\bar{y}$ lies on this geodesic and $\bar{x} \in S'_{\bar{x}_0}$. Let x be defined so that $[x_0, x]$ lifts $[\bar{x}_0, \bar{x}]$, and let y be the unique lift of $\bar{y}$ on this geodesic. By minimality, for all axes $h\text{Ax}(v) \in \mathcal{A}$, $D_{h\text{Ax}(v)}(x_0, x) \leq 0.51L(v)$. As $\bar{x} \in S'_{\bar{x}_0}$, there exists some lift x' of $\bar{x}$ such that x_0 and x' are minimal, and such that for all $h\text{Ax}(v) \in \mathcal{A}$, $D_{h\text{Ax}(v)}(x_0, x') \leq 0.39L(v)$.

From the triangle inequality, we conclude that for all $h\text{Ax}(v) \in \mathcal{A}$, $D_{h\text{Ax}(v)}(x, x') \leq 0.95L(v)$. From Lemma 4.1, we must have $x = x'$. Now, let $h\text{Ax}(v) \in \mathcal{A}$, so that $D_{h\text{Ax}(v)}(x_0, x) \leq 0.39L(v)$. From the Bounded Geodesic Image Theorem, the diameter of $\pi_{h\text{Ax}(v)}([x_0, x])$ is at most $0.395L(v)$. As $y \in [x_0, x]$, we conclude that $D_{h\text{Ax}(v)}(x_0, y) \leq 0.4L(v)$. $\square$

Lemma 8.13. *There exists a well-defined bi-Lipschitz map $\iota : S_{\bar{x}_0} \rightarrow X$.*

Proof. Given $\bar{x} \in S_{\bar{x}_0}$, define $\iota(\bar{x})$ to be a lift x of $\bar{x}$ such that x and x_0 are minimal. For any two such minimal lifts x_1, x_2 , we have by Lemma 8.12 that for any $hAx(v) \in \mathcal{A}$, $D_{hAx(v)}(x_0, x_i) \leq 0.4L(v)$ for $i \in \{1, 2\}$, so by the triangle inequality $D_{hAx(v)}(x_1, x_2) \leq 0.8L(v)$. But then Lemma 4.1 forces $x_1 = x_2$, and we conclude that ι is well-defined.

An analogous argument shows that for $\bar{x}, \bar{y} \in S_{\bar{x}_0}$, $x := \iota(\bar{x})$ and $y := \iota(\bar{y})$ are near-minimal. As a result, by Lemma 6.1,

$$d_X(x, y) \leq 10d_{X/N}(\bar{x}, \bar{y}).$$

On the other hand, certainly $d_{X/N}(\bar{x}, \bar{y}) \leq d_X(x, y)$. We obtain the desired result. $\square$

We note that for any $\bar{z} = g\bar{x}_0$, we have by equivariance of projections that $gS_{\bar{x}_0} = S_{\bar{z}}$. We therefore obtain the following corollary.

Corollary 8.14. *For any $\bar{z} \in X/N$, the set $S_{\bar{z}}$ is connected, quasi-convex, and quasi-isometric to a subset of X . As a result, $\partial S_{\bar{z}} := \Lambda(S_{\bar{z}})$ is closed in $\partial(X/N)$ and satisfies*

$$\dim(\partial S_{\bar{z}}) \leq \dim(\partial X).$$

We now introduce the main machinery. To do so, we first must define *stable struts*.

Definition 8.15. Let $\xi \in \partial(X/N)$, and let $\bar{r}$ be a $(1, 20L)$ -quasi-geodesic ray originating from $\bar{x}_0$ representing ξ . Let $hAx(v) \in \mathcal{A}/N$. We say that $hAx(v)$ is a *stable strut* for $\bar{r}$ if there exists a quantity $t_0 \geq 0$ satisfying the following two conditions.

1. For all $t \geq t_0$, $\overline{hAx(v)} \in \bar{S}(\bar{x}_0, \bar{r}(t))$
2. For all $t \geq t_0$, $d_{X/N}(\bar{r}(t), \overline{hAx(v)}) \geq 1000L$.

Definition 8.16. Define

$$S := \bigcup_{\bar{z} \in X/N} S_{\bar{z}}, \quad \partial S := \bigcup_{\bar{z} \in X/N} \partial S_{\bar{z}}.$$

Also define $\partial T = \partial S^C$. For each $\overline{hAx(v)} \in \mathcal{A}/N$, define $\partial T_{\overline{hAx(v)}}$ to be the set of all $\xi \in \partial T$ for which $\overline{hAx(v)}$ is a stable strut.

An illustration of the construction used in the following lemma is provided as Figure 10.

Lemma 8.17. *Let $\xi \in \partial(X/N)$, and assume that ξ is represented by some $(1, 20L)$ -quasi-geodesic ray $\bar{r}$ originating from $\bar{x}_0$ such that $\bar{r}$ has finitely many stable struts. Then there exists some $\bar{z} \in X/N$ such that $\xi \in \partial S_{\bar{z}}$.*

Proof. Note that key objects in this proof are illustrated in Figure 10.

By assumption, there exists some finite list $\overline{h_1 \text{Ax}(v_1)}, \dots, \overline{h_K \text{Ax}(v_K)}$ of stable struts for $\bar{r}$, and some t_0 such that for all $t \geq t_0$, we have $d_{X/N}(\overline{h_i \text{Ax}(v_i)}, \bar{r}(t)) \geq 1000L$ and $\overline{h_i \text{Ax}(v_i)} \in \overline{\mathcal{S}(\bar{x}_0, \bar{r}(t))}$ for $i \in \{1, \dots, K\}$. As a consequence of the fact that any $\overline{h_i \text{Ax}(v_i)}$ must lie within 50δ of $[\bar{x}_0, \bar{r}(t_0)]$ and each $\overline{h_i \text{Ax}(v_i)}$ has diameter at most $2L$, we have that

$$d_{X/N}(\bar{x}_0, \overline{h_i \text{Ax}(v_i)}) \leq d_{X/N}(\bar{x}_0, \bar{r}(t_0)) + 100L. \quad (2)$$

Now, let $\bar{y} = \bar{r}(t_0 + 10000L)$, and construct a geodesic $[\bar{x}_0, \bar{y}]$. Let $\bar{z}$ be a point on this geodesic such that $d_{X/N}(\bar{z}, \bar{y}) = 4000L$; certainly such a point exists. We will claim that $\xi \in \partial S_{\bar{z}}$.

Let $t \geq t_0 + 10000L$, and let $\bar{y}_t = \bar{r}(t)$. Construct a geodesic $[\bar{x}_0, \bar{y}_t]$, and let $\bar{z}_t$ be a point on this geodesic such that $d_{X/N}(\bar{z}_t, \bar{y}_t) = 2000L$. Our goal is to show that $\bar{z}_t \in S_{\bar{z}}$, and in fact $\bar{z}_t \in S'_{\bar{z}}$. Assume not, so there exist minimal respective lifts z_t and z of $\bar{z}_t$ and $\bar{z}$ and some axis $h\text{Ax}(v) \in \mathcal{A}$ with $D_{h\text{Ax}(v)}(z, z_t) > 0.39L(v)$. As $\bar{r}$ is a $(1, 20L)$ -quasi-geodesic, we may apply a standard thin-quadrilaterals argument and the fact that $\overline{h\text{Ax}(v)}$ has diameter at most $2L$ to obtain

$$d_{X/N}(\bar{x}_0, \overline{h\text{Ax}(v)}) \geq d_{X/N}(\bar{x}_0, \bar{r}(t_0)) + 2000L \quad (3)$$

We may now apply Lemma 5.10 to find a lift x_0 of $\bar{x}_0$ such that x_0 is near-minimal with respect to both z and z_t . From the triangle inequality, we must have that either $D_{h\text{Ax}(v)}(x_0, z)$ or $D_{h\text{Ax}(v)}(x_0, z_t)$ is at least $0.19L(v)$. Assume for now that $D_{h\text{Ax}(v)}(x_0, z) \geq 0.19L(v)$, so $\overline{h\text{Ax}(v)} \in \overline{\mathcal{S}(\bar{x}_0, \bar{z})}$. In particular, $\overline{h\text{Ax}(v)}$ must lie within 50δ of $[\bar{x}_0, \bar{z}]$, and hence

$$d_{X/N}(\bar{x}_0, \overline{h\text{Ax}(v)}) \leq d_{X/N}(\bar{x}_0, \bar{y}) - 1500L$$

As a consequence of Lemma 5.11 with $\theta = 0.15$, if z' is a separate lift of $\bar{z}$ such that $[x_0, z']$ lifts $[\bar{x}_0, \bar{z}] \subseteq [\bar{x}_0, \bar{y}]$, then there exists some lift $h'\text{Ax}(v')$ with $D_{h'\text{Ax}(v')}(x_0, z') \geq 0.15L(v')$. Note that x_0 and z' are necessarily minimal; we may furthermore extend $[x_0, z']$ to $[x_0, y']$ for some lift y' of $\bar{y}$ such that x_0 and y' are minimal. The Bounded Geodesic Image Theorem plus $[\bar{x}_0, \bar{z}] \subseteq [\bar{x}_0, \bar{y}]$ then implies $D_{h'\text{Ax}(v')}(x_0, y') \geq 0.14L(v')$. So $\overline{h'\text{Ax}(v')} \in \mathcal{S}_{0.14}(\bar{x}_0, \bar{y})$

Now, consider some $\bar{y}^* = \bar{r}(t^*)$ with $t^* \geq t_0 + 10000L$. As $\bar{r}$ is a $(1, 20L)$ -quasi-geodesic based at $\bar{x}_0$, we have that $[\bar{y}, \bar{y}^*]$ does not pass within $d_{X/N}(\bar{x}_0, \bar{y}) - 200L$ of $\bar{x}_0$. In particular, this geodesic does not pass within $1200L$ of $\overline{h\text{Ax}(v)}$. We may now find respective lifts y^* , x_0'' , and y'' of $\bar{y}^*$, $\bar{x}_0$, and $\bar{y}$ such that y^* and y'' are minimal, y^* and x_0'' are near-minimal, and y'' and x_0'' are near-minimal. From Lemma 5.11, there is some lift $h''\text{Ax}(v'')$ of $\overline{h\text{Ax}(v)}$ such that $D_{h''\text{Ax}(v'')}(x_0'', y'') \geq 0.14L(v)$. But as $[y'', y^*]$ can be chosen to not pass within $1200L$ of $h''\text{Ax}(v'')$ by the construction above, we may apply the Bounded Geodesic Image Theorem and conclude after projection that $\overline{h\text{Ax}(v)} \in \overline{\mathcal{S}(\bar{x}_0, \bar{y}^*)}$.

As t^* was chosen arbitrarily, we conclude that $\overline{hAx(v)}$ is a stable strut for $\bar{r}$. Combining Equations 2 and 3 now yields a contradiction.

We conclude that $\bar{z}_t \in S_{\bar{z}}$ for all $t \geq t_0 + 10000L$. As the sequence $\bar{z}_t$ $5000L$ -fellow-travels $\bar{r}$, it must serve as a defining sequence for ξ . We conclude that $\xi \in \Lambda(S_{\bar{z}})$. $\square$

Corollary 8.18. *If $\xi \in \partial T$, then ξ has infinitely many stable struts.*

We now investigate the $\partial T_{\overline{hAx(v)}}$ sets. An illustration of the construction used in the following lemma is provided as Figure 11.

Lemma 8.19. *For all $\overline{hAx(v)} \in \mathcal{A}/N$, the set $\partial T_{\overline{hAx(v)}}$ is clopen in the subspace topology on $\partial T \subseteq \partial(X/N)$. Furthermore, for any $\xi \in \partial T$, there exists a collection of $\partial T_{\overline{hAx(v)}}$ forming a basis at ξ in the subspace topology on $\partial(X/N)$.*

Proof. Note that key objects in the following construction are depicted in Figure 11.

If $\xi \in \partial T$, then there exists a $(1, 20L)$ -quasi-geodesic ray $\bar{r}$ representing ξ such that $\bar{r}$ has infinitely many distinct stable struts. Furthermore, let $\overline{h'Ax(v')}$ be a stable strut for $\bar{r}$, and consider some t such that for all $s \geq t$, $d_{X/N}(\bar{r}(s), \overline{hAx(v)}) \geq 1000L$. Certainly such a t exists. Then for all minimal respective lifts x_0 and $r(t)$, of $\bar{x}_0$ and $\bar{r}(t)$, there exists some lift $h'Ax(v')$ of $\overline{h'Ax(v')}$ such that, by Lemma 5.8 and the proof of Lemma 8.17, $D_{h'Ax(v')}(x_0, r(t)) \geq 0.099L(v')$. Therefore $[x_0, r(t)]$ must 100δ -fellow-travel $h'Ax(v')$ for distance at least $0.099L(v')$, and the same must hold after projection to the quotient. By the non-overlapping condition on axes, if $\bar{r}$ possesses infinitely many stable struts, then it must possess stable struts with arbitrary distance from $\bar{x}_0$.

Now, fix $\overline{hAx(v)} \in \mathcal{A}/N$, and let $\xi \in T_{\overline{hAx(v)}}$. From the preceding paragraph, we may find some $\overline{h'Ax(v')}$ such that $\overline{h'Ax(v')}$ is a stable strut for ξ , and such that

$$10d_{X/N}(\bar{x}_0, \overline{hAx(v)}) + 10000L \leq d_{X/N}(\bar{x}_0, \overline{h'Ax(v')})$$

There exists some t_0 such that $\overline{hAx(v)}, \overline{h'Ax(v')} \in \bar{\mathcal{S}}(\bar{x}_0, \bar{r}(t_0))$, and such that both $\overline{hAx(v)}$ and $\overline{h'Ax(v')}$ have distance at least $1000L$ from $\bar{r}(t_0)$. Applying the distance bound and Lemmas 5.11 and 5.8 to $[\bar{x}_0, \bar{r}(t_0)]$, we must have that $\overline{hAx(v)}$ comes before $\overline{h'Ax(v')}$ in $\bar{\mathcal{S}}(\bar{x}_0, \bar{r}(t_0))$, and so $\overline{hAx(v)} \in \bar{\mathcal{S}}(\bar{x}_0, \overline{h'Ax(v')})$.

Choose a constant

$$\mathcal{L} \geq d_{X/N}(\bar{x}_0, \overline{h'Ax(v')}) + 10000L$$

and consider the neighborhood

$$\{\chi \in \partial T : (\chi \cdot \xi)_{\bar{x}_0} \geq \mathcal{L}\}.$$

For any such χ , let $\bar{r}'$ be a $(1, 20L)$ -quasi-geodesic ray originating from $\bar{x}_0$ representing χ . Let $\bar{y}_\chi$ be a point on $\bar{r}'$ such that $d_{X/N}(\bar{x}_0, \bar{y}_\chi) \geq \mathcal{L} + 1000L$; note that there is certainly




Figure 10: A schematic of the construction used to prove Lemma 8.17, showing the relevant objects under projection to X/N . Note that the strut $\overline{hAx(v)}$ is proved to in fact not exist over the course of the proof.

some t_0 so that for all $t \geq t_0$, $\bar{r}'(t)$ satisfies this condition. Furthermore, $\bar{y}_\chi$ must satisfy $d_{X/N}(\overline{h'Ax(v')}, \bar{y}_\chi) \geq 1000L$. Likewise, choose a point $\bar{y}_\xi$ on $\bar{r}$ such that

$$d_{X/N}(\bar{x}_0, \bar{y}_\xi) = {}_{100L} \mathcal{L} + 1000L$$

From $(\chi \cdot \xi)_{\bar{x}_0} \geq \mathcal{L}$, the definition of $\mathcal{L}$, and a standard thin-triangles argument, any geodesic $[\bar{y}_\chi, \bar{y}_\xi]$ does not pass within $1000L$ of $\overline{h'Ax(v')}$.

We may then choose pairwise near-minimal respective lifts y_χ , y_ξ , and x_0 of $\bar{y}_\chi$, $\bar{y}_\xi$, and $\bar{x}_0$, so that y_χ and y_ξ are in fact minimal. Necessarily there exists some lift $h'Ax(v')$ of $\overline{h'Ax(v')}$ in $\mathcal{S}(x_0, y_\xi)$, and $D_{h'Ax(v')}(x_0, y_\xi) \geq 0.1L(v')$. There also exists some lift $hAx(v)$ of $\overline{hAx(v)}$ in $\mathcal{S}(x_0, y_\xi)$, and this axis satisfies

$$D_{hAx(v)}(x_0, y_\xi) \geq 0.1L(v), \quad D_{h'Ax(v')}(hAx(v), y_\xi) \geq 0.1L(v)$$

Here, the second inequality comes from the fact established above that $\overline{hAx(v)} \in \overline{\mathcal{S}(\bar{x}_0, \overline{h'Ax(v')})}$. As $h'Ax(v')$ has distance at least $1000L$ from any geodesic from y_ξ to y_χ by construction, the Bounded Geodesic Image Theorem implies that

$$D_{h'Ax(v')}(hAx(v), y_\chi) \geq 0.099L(v')$$

By forwarding, we therefore have

$$D_{hAx(v)}(x_0, y_\chi) \geq 0.1L(v)$$

and as x_0 and y_χ are near-minimal, $\overline{hAx(v)} \in \overline{\mathcal{S}(\bar{x}_0, \bar{y}_\chi)}$. From this analysis, we may conclude that $\overline{hAx(v)}$ is a stable strut for χ , and hence $\chi \in \partial T_{\overline{hAx(v)}}$. So

$$\{\chi \in \partial T : (\chi \cdot \xi)_{\bar{x}_0} \geq \mathcal{L}\} \subseteq \partial T_{\overline{hAx(v)}}.$$

Next, assume instead that $\xi \notin T_{\overline{hAx(v)}}$. We choose $\overline{h'Ax(v')}$ and $\mathcal{L}$ as before. Let $\bar{y}_\xi$ be a point on $\bar{r}$ such that $\overline{h'Ax(v')} \in \overline{\mathcal{S}(\bar{x}_0, \bar{y})}$ and $d_{X/N}(\overline{h'Ax(v')}, \bar{y}_\xi) \geq 10000L$. Without loss of generality, we may also assume that $d_{X/N}(\bar{x}_0, \bar{y}_\xi) \geq \mathcal{L} + 1000L$. If $\overline{hAx(v)} \in \overline{\mathcal{S}(\bar{x}_0, \overline{h'Ax(v')})}$, we would have $\overline{hAx(v)} \in \overline{\mathcal{S}(\bar{x}_0, \bar{y}_\xi)}$. Then setting $\chi = \xi$ and rerunning the preceding argument would yield that $\overline{hAx(v)}$ is a stable strut for ξ , contradicting $\xi \notin \partial T_{\overline{hAx(v)}}$. We conclude that for any pairwise minimal lifts x_0 and $h'Ax(v')$ of $\bar{x}_0$ and $\overline{h'Ax(v')}$ and any lift $hAx(v)$ of $\overline{hAx(v)}$, we have $D_{hAx(v)}(x_0, h'Ax(v')) < 0.1L(v)$.

Consider the neighborhood

$$\{\chi \in \partial T : (\chi \cdot \xi)_{\bar{x}_0} \geq \mathcal{L}\}$$

and assume there exists some χ lying in this neighborhood so that $\chi \in \partial T_{\overline{hAx(v)}}$. We may choose a $(1, 20L)$ -quasi-geodesic ray $\bar{r}'$ representing χ ; for all y_χ on $\bar{r}'$ sufficiently far from

$\bar{x}_0$, we have that $\overline{hAx(v)} \in \bar{S}(\bar{x}_0, \bar{y}_\chi)$ and $d_{X/N}(\bar{x}_0, \bar{y}_\xi) \geq \mathcal{L} + 1000L$. As before, by choice of $\mathcal{L}$ and $(\chi \cdot \xi)_{\bar{x}_0} \geq \mathcal{L}$, any geodesic in X/N from $\bar{y}_\xi$ to $\bar{y}_\chi$ must not pass within $1000L$ of $\overline{h'Ax(v')}$

By essentially the same forwarding argument as in the previous case, plus the upper bound on $D_{hAx(v)}(x_0, h'Ax(v'))$ given above, we conclude that

$$\{\chi \in \partial(X/N) : (\chi \cdot \xi)_{\bar{x}_0} \geq \mathcal{L}\} \subseteq \partial T_{hAx(v)}^C$$

As sets of the form

$$\{\chi \in \partial(X/N) : (\chi \cdot \xi)_{\bar{x}_0} \geq R\}$$

for $R \geq 0$ form a basis neighborhood for ξ in the visual topology on $\partial(X/N)$, we conclude that $\partial T_{hAx(v)}$ is both closed and open.

Lastly, for any $\overline{hAx(v)} \in \mathcal{A}/N$ and for any $\xi, \chi \in \partial T_{hAx(v)}$, it is straightforward to verify from Lemma 5.8 and the stable strut condition that any $(1, 20L)$ -quasi-geodesic rays representing ξ or χ must both pass within, say, $100L$ of $\overline{hAx(v)}$. As $\overline{hAx(v)}$ has diameter at most $2L$, we may conclude that

$$(\xi \cdot \chi)_{\bar{x}_0} \geq d_{X/N}(\bar{x}_0, \overline{hAx(v)}) - 1000L$$

For any ξ , we may find stable struts for ξ with arbitrary distance from $\bar{x}_0$. Constructing sets $\partial T_{hAx(v)}$ from these stable struts therefore gives a neighborhood basis for ξ . $\square$

As every point of ∂T has a neighborhood basis of clopen sets of the form $\partial T_{hAx(v)}$.

Corollary 8.20. *The space $\partial T \subseteq \partial(X/N)$ equipped with the subspace topology has topological dimension 0.*

At last, we may prove our main result.

Proof of Lemma 8.3. Recall the definitions

$$\partial S = \bigcup_{\bar{z} \in X/N} \Lambda(S_{\bar{z}}), \quad \partial T = \partial S^C.$$

As the action of G on X/N is transitive and G is finitely generated, the union is countable. From Lemma 8.4 and Corollary 8.14, we obtain $\dim(\partial S) \leq \dim(\partial X)$, with the subspace topology on ∂S . As $\dim(\partial T) = 0$ by Corollary 8.20, we may apply Lemma 8.5 to conclude that as $\partial(X/N) = \partial S \cup \partial T$,

$$\dim(\partial(X/N)) \leq \dim(\partial S) + \dim(\partial T) + 1 = \dim(\partial S) + 1.$$

Lastly, by Theorem 7.8, $\partial(X/N)$ is linearly connected, and hence has dimension at least 1. $\square$




Figure 11: A schematic of the construction used to prove Lemma 8.19, showing the relevant objects under projection to X/N . The key idea is that the struts $\overline{h'_1 \text{Ax}(v'_1)}$ and $\overline{h'_2 \text{Ax}(v'_2)}$ ‘lock in’ projection angles for (respectively) boundary points near ξ_1 and ξ_2 to $\overline{h \text{Ax}(v)}$. Here, $\xi_1 \in \partial T_{\overline{h \text{Ax}(v)}}$, whereas $\xi_2 \notin \partial T_{\overline{h \text{Ax}(v)}}$

9 Boundary Quotients and the Nöbeling Curve

9.1 The Nöbeling Curve

We now begin the proof of Theorem 1.3, motivated largely by the results of [HP11]. First, we provide a brief introduction to the Nöbeling Curve.

Definition 9.1. The *Nöbeling Curve* is the topological space defined by removing from $\mathbb{R}^3$ all points with at least two rational coordinates and taking the subspace topology.

It is reasonable to think of the Nöbeling curve as a non-compact analogue of the Menger curve (also known as the Menger sponge).

The Nöbeling curve admits a characterization theorem. It may also be defined via a universal property, although we will not state this definition. Before stating the theorem, we must define a somewhat peculiar property.

Definition 9.2. A topological space X satisfies the *locally finite 1-disks property* if the following conditions are satisfied. For any family of continuous maps $f_m : [0, 1] \rightarrow X$ and any open cover $\mathcal{U}$ of X , there is a family of continuous maps $g_m : [0, 1] \rightarrow X$ such that:

1. For each $x \in X$, there is an open neighborhood U of x (not necessarily in $\mathcal{U}$) such that the image of g_n does not intersect U for all sufficiently large n .
2. For each $t \in [0, 1]$ and $n \in \mathbb{N}$, there is an open set $U \in \mathcal{U}$ such that $f_n(t)$ and $g_n(t)$ both lie in U .

For an exposition on why the Nöbeling curve possesses the locally finite 1-disks property, see Prop. 6.1 in [HP11]. The following result was proven by Andrej Andrzej Nagórko based on results in [KLT97], but is stated in [HP11].

Theorem 9.3. *A topological space U is homeomorphic to the Nöbeling curve if and only if the following properties are satisfied.*

1. *U is Polish (separable and complete metrizable).*
2. *U is connected and locally path-connected.*
3. *U has topological dimension 1.*
4. *U possesses the locally finite 1-disks property.*

Theorem 1.3 will require additional hypotheses, which we now cover. The first hypothesis is a strengthening of Hypothesis 3.1.

9.2 Hypotheses

We now introduce our main hypothesis for the Nöbeling curve result. Note that relative hierarchical hyperbolicity is a strictly weaker condition than hierarchical hyperbolicity.

Note that any HHG G comes with an action $G \curvearrowright \mathcal{X}$ on a *hierarchically hyperbolic space* $\mathcal{X}$, a set $\mathcal{D}$ of domains with a unique maximal domain X , and an action $G \curvearrowright \mathcal{D}$. For any $U \in \mathcal{D} \setminus \{X\}$, we have a projection map $\rho_U^X : X \rightarrow U$ and a coarse point $\rho_X^U \subseteq X$. There is also an orthogonality relationship $\perp$ on elements of $\mathcal{D}$. Furthermore, for any $U \in \mathcal{D}$, there is a projection map $\pi_U : \mathcal{X} \rightarrow U$; these maps satisfy a coarse functoriality property.

Hypothesis 9.4. *Let G be a finitely generated relatively hierarchically hyperbolic group (RHHG) such that the RHHG structure of G is realized by an action $G \curvearrowright \mathcal{X}$, where $\mathcal{X}$ is a relatively hierarchically hyperbolic space. Let $\mathcal{D}$ be the collection of domains for $\mathcal{X}$, and recall that by definition there is an action $G \curvearrowright \mathcal{D}$. Let E be the maximum across all constants appearing in the RHHS/RHHG axioms for G and $\mathcal{X}$. Let X be the top level domain of $\mathcal{X}$. By definition, X is a Gromov hyperbolic space, and there is an action $G \curvearrowright X$ of G on X by hyperbolic isometries. We let $w_1, \dots, w_k$ be elements of G . We fix an arbitrary basepoint $x_0 \in X$, and consider the normal subgroup $N = \langle\langle w_1, \dots, w_k \rangle\rangle$. We assume the following:*

1. *The space $\mathcal{X}$ is the Cayley graph of G with respect to some finite generating set, and is equipped with the simplicial metric.*
2. *The space X is a simplicial graph, and its Gromov hyperbolicity constant δ is at least 1000.*
3. *The action of G on X is vertex-transitive and non-elementary.*
4. *The maximal finite normal subgroup of G is trivial.*
5. *For each non-maximal $U \in \mathcal{D}$, ρ_X^U is a point in X . Furthermore, for $g \in G$, if $U' = gU$, then $\rho_X^{U'} = g\rho_X^U$.*
6. *There exists $M \geq 10^{20}$ such that*

$$(w_1, \dots, w_k) \in \mathcal{E}_{10^{20}M, M}$$

where $\mathcal{E}_{10^{20}M, M}$ is defined as in Definition 2.25.

Given these hypotheses, we set $L = \max_i L(w_i)$ and $\mathcal{M} := 50 \min_i L(w_i)/M$.

Remark 9.5. It is straightforward to extend the reductions applied to Hypothesis 3.1 to this expanded setup; see Appendix A of [Abb+26] for the needed facts about quotients of $\mathcal{X}$ by finite normal subgroups of G . As before, the other claims are standard.

We note that the action of G on X is acylindrical, and hence WPD, by Theorem K of [BHS17] (the proof is only for HHGs, but the argument works for RHHGs as well), and the action is likewise action path-connected by Lemma 2.29. Vertex-transitivity and conditions four and five may likewise be assumed without loss of generality. Possibly replacing E with a larger constant, we can assume that E also serves as the action path-connectedness constant.

We will use ‘subdomain’ to refer to non-maximal domains of $\mathcal{X}$; note that this is not standard terminology.

We make the following definition, which is the main requirement for showing that $\partial(X/N)$ has the locally finite 1-disks property.

Definition 9.6. Let G and $\mathcal{X}$ be as in Hypothesis 9.4. Recall We say that $\mathcal{X}$ has the *large dense domains* property if there exists a subset $\mathcal{D}_0 \subseteq \mathcal{D} \setminus \{X\}$ such that the following conditions are satisfied for some constant H :

1. The subset $\mathcal{D}_0$ is closed under the G -action on $\mathcal{D}$.
2. Every $U \in \mathcal{D}_0$ contains an quasi-isometrically embedded copy of $\mathbb{R}$.
3. Every point $x \in X$ lies within distance H of some $U \in \mathcal{D}_0$.
4. For every geodesic $[x, y]$ in X , there is a sequence $U_1, \dots, U_k$ of domains in $\mathcal{D}_0$ so that $\rho_X^{U_1}$ lies within H of x , $\rho_X^{U_k}$ lies within H of y , every $\rho_X^{U_i}$ lies within H of $[x, y]$, and $U_i \perp U_{i+1}$ and $d_X(\rho_X^{U_i}, \rho_X^{U_{i+1}}) \leq 10H$ for all U_i with $1 \leq i \leq k-1$.

We make the following hypothesis.

Hypothesis 9.7. Let G , $\mathcal{X}$, and X be as in Hypothesis 9.4. We additionally assume that G , $\mathcal{X}$, and X satisfy the large dense domains property.

If needed, we may replace E with a larger constant so that $E \geq H$.

Lastly, we make the following hypothesis on X . Note that quasi-trees necessarily have boundaries with dimension 0.

Hypothesis 9.8. Let G , $\mathcal{X}$, and X be as in Hypothesis 9.4. We additionally assume that X is a quasi-tree.

10 Locally Finite 1-Disks

In this section, we prove that under Hypotheses 9.4 and 9.7, $\partial(X/N)$ satisfies the locally finite 1-disks property. Before we begin, recall the triangle inequality for Gromov products: for any ξ_1, ξ_2, ξ_3 that are individually either points in X/N or points in $\partial(X/N)$ and any $\bar{x} \in X/N$, we by Lemma 4.5 that

$$(\xi_1 \cdot \xi_3)_{\bar{x}_0} \geq \min((\xi_1 \cdot \xi_2)_{\bar{x}_0}, (\xi_2 \cdot \xi_3)_{\bar{x}_0}) - 20L$$

10.1 Subdomain Projections in the Quotient

Following Abbott-Berlyne-Mangioni-Ng-Rasmussen in [Abb+26], we could define a full RHHS structure on $\mathcal{X}/N$, and obtain subdomain projections accordingly. However, this is more work than is necessary in our case, and we will instead define and prove only what is needed to show the locally finite 1-disks property.

Recall that for reach

Definition 10.1. Define $\overline{\mathcal{D}}$ as the set-theoretic quotient $\mathcal{D}/N$ under the action $G \curvearrowright \mathcal{D}$, except that X is replaced by X/N . Define $\overline{\mathcal{D}}_0 := \mathcal{D}_0/N \subseteq \mathcal{D}/N$ accordingly. For each $\overline{U} \in \overline{\mathcal{D}} \setminus \{X/N\}$, we define $\overline{U}$ as an isometric copy of any $U \in \mathcal{D}$ lifting $\overline{U}$. Given $\overline{U} \in \overline{\mathcal{D}}$, define $\rho_{\overline{X}}^{\overline{U}}$ as the projection of ρ_X^U to $\overline{X}$ for any lift U of $\overline{U}$. Note that this point is well-defined by Hypothesis 9.4.

Recall that by definition, if $U' = gU$ for distinct $U, U' \in \mathcal{D}$, then the map from U to U' induced by g is an isometry. Furthermore, $\rho_X^{gU} = g\rho_X^U$, and as all $g \in N$ have minimal translation length at least $0.98 \min_i L(v_i)$ by Lemma 4.1, we may conclude N acts freely on $\mathcal{D} \setminus \{X\}$. Therefore, if $U' = gU$ for $g \in N$, the induced isometry between U' and U is unique. As a result, we may define a projection map

$$\pi_N : \bigsqcup_{U \text{ lifts } \overline{U}} U \rightarrow \overline{U}$$

by identifying $\overline{U}$ with a fixed lift U_0 of $\overline{U}$, mapping all other lifts U to U_0 via the unique maps described above, and taking a union of images.

Recall that for each $U \in \mathcal{D} \setminus \{X\}$, there exists a ‘projection’ map $\rho_U^X : X \rightarrow U$.

Definition 10.2. Given $\overline{x} \in X/N$ and $\overline{U} \in \overline{\mathcal{D}} \setminus \{X/N\}$, define

$$\rho_{\overline{U}}^{\overline{X}}(\overline{x}) = \left(\bigsqcup_{\substack{x, U: \\ x \text{ lifts } \overline{x} \\ U \text{ lifts } \overline{U} \\ x, U \text{ are 0.85-near-minimal}}} \rho_U^X(x) \right) / N$$

where the quotient by N is defined via the projection map π_N described in Definition 10.1.

Before we go on, recall the Bounded Geodesic Image Axiom for RHHSes.

Lemma 10.3. *Let G , $\mathcal{X}$, and X be as in Hypothesis 9.4. Let x and y be points of X such that $[x, y]$ does not intersect $\mathcal{N}_E(\rho_X^U)$ for some subdomain $U \in \mathcal{D} \setminus \{X\}$. Then $d_U(\rho_U^X(x), \rho_U^X(y)) \leq E$.*

We now show projections have bounded diameter. A helper lemma is required; this lemma is essential Remark 4.17 in [Abb+26].

Lemma 10.4. *There exists a constant E' depending only on E and the action $G \curvearrowright X$ with the following property. Let $hAx(v)$ be an axis, let $g = hvh^{-1}$, and let x be a point lying on $hAx(v)$. Then for any $U \in \mathcal{D}$, $d_U(\rho_U^X(x), \rho_U^X(gx)) \leq E'$.*

Proof. The existence of $\hat{x}$ follows from transitivity of the action of G on X . We note that $\langle hvh^{-1} \rangle$ acts geometrically on the quasi-convex subset $hAx(v)$, which implies the existence of constants (λ, c) such that the restriction of the map $\pi_X : \mathcal{X} \rightarrow X$ (which is the orbit map for G) to $\langle hvh^{-1} \rangle$ is a (λ, c) -quasi-isometry. It now follows from the Bounded Geodesic Image Axiom that the image $\rho_U^X(hAx(v))$ has diameter bounded by some function $D(\lambda, c)$ of λ and c . As G acts on the set $\mathcal{A}$ of axes with finitely many orbits, λ and c are uniformly bounded across all axes. We therefore obtain a bound on the diameter of $\rho_U^X(hAx(v))$ independent of $hAx(v) \in \mathcal{A}$. $\square$

We will assume without loss of generality that $E' \leq E$, which may require replacing E with a larger constant.

Lemma 10.5. *With the notation in Definition 10.2, the set $\rho_{\overline{U}}^{\overline{X}}(\overline{x})$ has diameter at most $O(E)$.*

Proof. Consider pairs x_1, U_1 and x_2, U_2 , both lifting the pair $\overline{x}, \overline{U}$, and such that x_1 and U_1 are 0.85-near-minimal and x_2 and U_2 are 0.85-near-minimal. By definition of $\rho_{\overline{U}}^{\overline{X}}(\overline{x})$, we may translate by N and thus assume without loss of generality that $U_1 = U_2$. Call this subdomain U . If $x_1 = x_2$, we are done. Otherwise, we may apply Lemma 4.1 and conclude that there is some axis $hAx(v)$ such that $d_{hAx(v)}(x_1, x_2) \geq 0.99L(v)$.

Now, if $d_{hAx(v)}(x_i, \rho_X^U) \leq 0.13L(v)$ for $i \in \{1, 2\}$, we may conclude by the triangle inequality that either x_1 or x_2 is not 0.85-near-minimal with U . As a consequence, defining y_1 as an arbitrary point in $\pi_{hAx(v)}(x_1)$, any geodesic from x_1 to ρ_X^U must first pass within $O(\delta)$ of y_1 and then travel along $hAx(v)$ for distance at least $0.1L(v)$. We may define y_2 analogously. By Lemma 10.3, as long as $L(v)$ is sufficiently large relative to E (which we may assume), $d_U(\rho_U^X(x_i), \rho_U^X(y_i)) \leq E$ for $i \in \{1, 2\}$.

It is now straightforward to verify that there exists an element $g = hv^m h^{-1}$ with $m \in \{-1, 0, 1\}$ such that $[gy_1, y_2]$ does not pass within E of ρ_X^U , and so $d_U(\rho_U^X(gy_1), \rho_U^X(y_2)) \leq E$. Furthermore, by Lemma 10.4, $d_U(\rho_U^X(y_1), \rho_U^X(gy_1)) \leq E$. We may now apply the Bounded Geodesic Image axiom a final time and conclude by the triangle inequality that

$$d_U(\rho_U^X(x_1), \rho_U^X(x_2)) \leq O(E). \square$$

Lemma 10.6. *Let $\overline{x}, \overline{y} \in X/N$, and consider $\overline{U} \in \overline{\mathcal{D}}$. Assume that $[\overline{x}, \overline{y}]$ does not pass within E of $\rho_{\overline{X}}^{\overline{U}}$. Then $d_{\overline{U}}(\rho_{\overline{U}}^{\overline{X}}(\overline{x}), \rho_{\overline{U}}^{\overline{X}}(\overline{y})) \leq O(E)$.*

Proof. Let x, U_1 be near-minimal lifts of $\bar{x}, \bar{U}$, and let y, U_2 be near-minimal lifts of $\bar{y}, \bar{U}$. By equivariance and a minor modification of Lemma 5.10, we may assume that $U_1 = U_2$ (call this subdomain U) and that x and y are minimal. By Lemma 10.5, it suffices to prove that $d_U(\rho_U^X(x), \rho_U^X(y)) \leq E$. But this bound now holds from our construction and the Bounded Geodesic Image axiom for X . $\square$

Definition 10.7. Given a quasi-geodesic ray $\bar{r}$ in X/N and a point $\bar{y}$ on $\bar{r}$, define $\bar{r}_{\bar{y}}$ to be the ‘tail’ of $\bar{r}$ obtained by removing all points of $\bar{r}$ before $\bar{y}$.

Definition 10.8. Fix $\xi \in \partial(X/N)$ and $\bar{U} \in \bar{\mathcal{D}} \setminus \{X/N\}$. Let R be the set of all $(1, 20L)$ -quasi-geodesic rays $\bar{r}$ in X/N originating from $\bar{x}_0$ and representing ξ . Then

$$\rho_{\bar{U}}^{\partial(X/N)}(\xi) = \bigcup_{\bar{r} \in R} \bigcup_{\substack{\bar{y} \in \bar{r} \\ d_{X/N}(\bar{r}_{\bar{y}}, \rho_{\bar{U}}^{\bar{X}}) \geq 10^{10}L}} \rho_{\bar{U}}^{\bar{X}}(\bar{y})$$

The following lemma follows directly from Lemmas 10.5 and 10.6.

Lemma 10.9. *For any $\xi \in \partial(X/N)$ and any $\bar{U} \in \bar{\mathcal{D}} \setminus \{X/N\}$, the set $\rho_{\bar{U}}^{\partial(X/N)}(\xi)$ has diameter at most $O(E)$.*

Definition 10.10. Given $\bar{U}, \bar{V} \in \bar{\mathcal{D}} \setminus \{X/N\}$, we write $\bar{U} \perp \bar{V}$ if there exist near-minimal respective lifts U and V of $\bar{U}$ and $\bar{V}$ with $U \perp V$.

We record a useful consequence of partial realization.

Lemma 10.11. *Fix $\bar{U} \in \bar{\mathcal{D}} \setminus \{X/N\}$. Then the map $\rho_{\bar{U}}^{\partial(X/N)} \rightarrow \bar{U}$ is $O(E)$ -coarsely surjective.*

Proof. Fix $\bar{u} \in \bar{U}$, and let U be an arbitrary lift of $\bar{U}$. By partial realization, there exists a point $\hat{u} \in X$ and a point $u' \in U$ projecting to $\bar{u}' \in \bar{U}$ such that u' is contained in the projection of $\hat{u}$ to U and $d_{\bar{U}}(\bar{u}', \bar{u}) \leq E$. Let x_u be a point in the projection of $\hat{u}$ to X , and let $\bar{x}_u$ be its image in X/N . Possibly after adjusting $\hat{u}$ and u' in a way that moves u' by distance at most E , we may assume without loss of generality that $d_X(x_u, \rho_X^U) \leq E$, and that as a consequence x_u and U are near-minimal.

Say $x_u = hx_0$. As in the proof of Lemma 7.3, we will choose a group element $c \in G$ with the following properties:

1. There exists a $c' \in G$ and a v_i so that $[x_0, cx_0] \subseteq \mathcal{N}_{20\delta}((c')^{-1}\text{Ax}(v_i))$.
2. We have $d_X(x_0, cx_0) =_{100\delta} \frac{1}{3}L_i$.
3. We have $d_G(e, c) = O(n)$.

We then take the sequence $z_j = hc^j x_0$, so $z_0 = x_u$. By essentially the same proof as in Lemma 7.4, the sequence $\bar{z}_j$ defines a boundary point $\xi \in \partial(X/N)$, and all z_j are near-minimal with respect to U . Furthermore, we may concatenate geodesic segments $[\bar{z}_j, \bar{z}_{j+1}]$ to obtain a $(1, 1000\mathcal{M})$ -quasi-geodesic ray $\bar{s}$. We may subdivide this ray into at most three sections $\bar{s}_1$, $\bar{s}_2$, and $\bar{s}_3$, each section made of some number $[z_j, z_{j+1}]$ segments, and satisfying the following properties. We require that $\bar{s}_1$ and $\bar{s}_3$ do not pass within $10L$ of $\rho_X^{\bar{U}}$, and that $\bar{s}_2$ consists of at most 40 segments. Let z_{j_1} be the terminal point of $\bar{s}_1$, let z_{j_2} be the terminal point of $\bar{s}_2$, and let $\bar{z}_{j_3}$ be some arbitrary point on $\bar{s}_3$ with distance at least $10^{20}L$ from $\rho_X^{\bar{U}}$. By the word-length bound above, the coarse Lipschitz projections property, and Lemma 10.6, we have

$$d_{\bar{U}}(\rho_X^{\bar{U}}(\bar{z}_0), \rho_X^{\bar{U}}(\bar{z}_{j_1})) \leq O(E), \quad d_{\bar{U}}(\rho_X^{\bar{U}}(\bar{z}_{j_1}), \rho_X^{\bar{U}}(\bar{z}_{j_2})) \leq O(E),$$

$$d_{\bar{U}}(\rho_X^{\bar{U}}(\bar{z}_{j_2}), \rho_X^{\bar{U}}(\bar{z}_{j_3})) \leq O(E)$$

Let $\bar{r}$ be a $(1, 20L)$ -quasi-geodesic ray in X/N originating from $\bar{x}_0$ and representing ξ . By the Morse Lemma, $\bar{r}$ fellow-travels the sequence $\bar{z}_j$ for sufficiently large j . We may now choose a sufficiently large j_3 and conclude by the triangle inequality. $\square$

We now recall a specific case of the Partial Realization Axiom for HHses.

Lemma 10.12. *Let $U, V \in \mathcal{D} \setminus \{X\}$ with $U \perp V$. Then the map*

$$\pi_U \times \pi_V : \mathcal{X} \rightarrow U \times V$$

is E -coarsely surjective.

The following lemma now follows immediately from the proof of Lemma 10.11 and the Partial Realization Axiom.

Lemma 10.13. *Fix $\bar{U}, \bar{V} \in \overline{\mathcal{D}} \setminus \{X/N\}$, and assume that $\bar{U} \perp \bar{V}$. Then the map*

$$\rho_{\bar{U}}^{\partial(X/N)} \times \rho_{\bar{V}}^{\partial(X/N)} : \partial(X/N) \rightarrow \bar{U} \times \bar{V}$$

is $O(n)$ -coarsely surjective.

We now state an extension of the Bounded Geodesic Image Axiom that applies to boundary points in the quotient. This lemma follows directly from unraveling definitions and applying Lemma 10.6.

Lemma 10.14. *There exists a constant $J = O(L)$ with the following property. Let $\bar{U} \in \overline{\mathcal{D}} \setminus \{X/N\}$ and $\xi_1, \xi_2 \in \partial(X/N)$. If $d_{\bar{U}}(\rho_{\bar{U}}^{\partial(X/N)}(\xi_1), \rho_{\bar{U}}^{\bar{X}}(\bar{x}_0))$ is at least J , then any $(1, 20L)$ -quasi-geodesic ray originating at $\bar{x}_0$ representing ξ passes within J of $\rho_X^{\bar{U}}$. Furthermore, if $d_{\bar{U}}(\rho_{\bar{U}}^{\partial(X/N)}(\xi_1), \rho_{\bar{U}}^{\partial(X/N)}(\xi_2))$ is at least J , then $(\xi_1 \cdot \xi_2)_{\rho_X^{\bar{U}}} \leq E$.*

We now reach the most technical claim, which will require that we rerun the construction from Lemma 7.1 with a few modifications.

Lemma 10.15. *For any constant $\Psi_1 > 0$, there exists a constant $\Psi_2 > 0$ with the following property. Let $\xi_1, \xi_2 \in \partial(X/N)$ and $\bar{U} \in \bar{\mathcal{D}} \setminus \{X/N\}$, and assume that $d_{\bar{U}}(\rho_{\bar{U}}^{\partial(X/N)}(\xi_1), \rho_{\bar{U}}^{\partial(X/N)}(\xi_2)) \leq \Psi_1$. Then there exists a path $\gamma : [0, 1] \rightarrow \partial(X/N)$ from ξ_1 to ξ_2 such that for all $t \in [0, 1]$, $d_{\bar{U}}(\rho_{\bar{U}}^{\partial(X/N)}(\gamma(t)), \rho_{\bar{U}}^{\partial(X/N)}(\xi_1)) \leq \Psi_2$.*

In order to prove this lemma, we make the following claim:

Lemma 10.16. *Assume Hypothesis 9.4. Fix $\xi_1, \xi_2 \in \partial(X/N)$ with $(\xi_1 \cdot \xi_2)_{\bar{x}_0} = D$ and $(\xi_1 \cdot \xi_2)_{\rho_{\bar{X}}} = D'$, and a constant F . Also fix $\bar{U} \in \bar{\mathcal{D}} \setminus \{X/N\}$, and assume that $d_{\bar{U}}(\rho_{\bar{U}}^{\partial(X/N)}(\xi_1), \rho_{\bar{U}}^{\partial(X/N)}(\xi_2)) \leq \Psi_1$. Then for M sufficiently large in terms of F , the following property holds: there exists a finite sequence of boundary points $\xi_1 = \chi_1, \chi_2, \dots, \chi_Q = \xi_2$ such that for $1 \leq i \leq Q-1$, $(\chi_i \cdot \chi_{i+1})_{\bar{x}_0} \geq D + \frac{1}{2}F$, and for all $2 \leq i \leq Q-1$, and $j \in \{1, Q\}$, $(\chi_i \cdot \chi_j)_{\bar{x}_0} \geq D - 100L$. Furthermore, for all χ_i with $1 \leq i \leq Q-1$ and all χ_j with $j \in \{1, Q\}$, we have $(\chi_i \cdot \chi_{i+1})_{\rho_{\bar{X}}} \geq D' + \frac{1}{2}F$ as well as $d_{\bar{U}}(\rho_{\bar{U}}^{\partial(X/N)}(\chi_i), \rho_{\bar{U}}^{\partial(X/N)}(\chi_j)) \leq \Psi_1 + O(L)$.*

Proof. The proof of this lemma will be substantially derived from the proof of Lemma 7.1, so we will gloss over the aspects that remain unchanged. We will also borrow key notation from this proof. As before, we let $\bar{r}_1, \bar{r}_2 : \mathbb{Z}_{\geq 0} \rightarrow X/N$ be $(1, 20L)$ -quasi-geodesic rays originating at $\bar{x}_0$ and respectively representing ξ_1 and ξ_2 . We choose points $\bar{x}_1$ and $\bar{x}_2$ on these rays such that $d_{X/N}(\bar{x}_0, \bar{x}_1) \geq \mathcal{Z}$ and $(\bar{x}_1, \bar{x}_2)_{\bar{x}_0} =_{10} D$. We now also require that $d_{X/N}(\bar{x}_i, \rho_{\bar{X}}^{\bar{U}}) \geq \mathcal{Z} + 10L$. Here, $\mathcal{Z}$ is some large multiple of L , say $10^{12}L$; it is straightforward that $\bar{x}_i$ exist satisfying this additional bound.

We let x_0, x_1, x_2 be pairwise near-minimal respective lifts of $\bar{x}_0, \bar{x}_1, \bar{x}_2$ to X , chosen so that x_1 and x_2 are minimal. Let γ be a geodesic from x_1 to x_2 , and set $\bar{\gamma} = \pi_{X/N}(\gamma)$. We now use Theorem 2.28 to construct a hierarchy path from x_1 to x_2 (this construction replaces the use of action path-connectedness). We obtain a sequence of group elements $a_1, \dots, a_k \in G$ such that $a_1 x_0 = x_1$, $a_k x_0 = x_2$, such that $d_G(a_{i-1}, a_i) \leq W$ for all $2 \leq i \leq k$, and such that $x_1 = a_1 x_0, a_2 x_0, \dots, a_k x_0 = x_2$ is an unparameterized (E, E) -quasigeodesic. As before, we define $\bar{w}_i := a_i \bar{x}_0$. By Theorem 2.28, the Bounded Geodesic Image Axiom, and our choice of $\mathcal{Z}$, the assumption that

$$d_{\bar{U}}(\rho_{\bar{U}}^{\partial(X/N)}(\xi_1), \rho_{\bar{U}}^{\partial(X/N)}(\xi_2)) \leq \Psi_1$$

implies that

$$d_{\bar{U}}(\rho_{\bar{U}}^{\bar{X}}(\bar{x}_1), \rho_{\bar{U}}^{\bar{X}}(\bar{w}_i)) \leq \Psi_1 + O(E^3).$$

Recall that in Lemma 7.2 we constructed points $u_i := a_i s_i x_0$ avoiding certain ‘enemy prongs’; using the notation of Lemma 6.2, two of these prongs were $p_1(w_i, \bar{x}_0)$ and $p_2(w_i, \bar{x}_0)$. We add $p_1(w_i, \rho_{\bar{X}}^{\bar{U}})$ and $p_2(w_i, \rho_{\bar{X}}^{\bar{U}})$ to the list of enemy prongs we wish to avoid

with u_i . Recall that $d_G(e, s_i) \leq E_M$ by the Schottky set construction. As a consequence of the coarse Lipschitz property of projections from $\mathcal{X}$ and the functoriality axiom for projections,

$$d_{\overline{U}}(\rho_{\overline{U}}^{\overline{X}}(\overline{w}_i), \rho_{\overline{U}}^{\overline{X}}(\overline{u}_i)) \leq O(EE_M) + O(E) = O(B)$$

Here, B is the constant used in Lemma 7.1.

While the statement was not needed before, it may be derived from the proof of the second case of Lemma 7.7 that for any $(1, 20L)$ -quasi-geodesic ray ρ representing some $\chi_{i,j}$ and any point $\overline{s}$ sufficiently far along on ρ , any geodesic $[\overline{s}, \overline{x}_0]$ passes through some point $\overline{y}$ lying within distance 100δ of either $\overline{u}_i$ or $\overline{u}_{i+1}$, and furthermore the subgeodesic $[\overline{y}, \overline{x}_0]$ passes within $1200B$ of $\overline{w}_i$. We apply essentially the same argument in our new setup, using both the enemy prongs constructed from $\rho_{\overline{X}}^{\overline{U}}$ and the enemy prongs constructed from $\overline{x}_0$. We obtain that any geodesic $[\overline{s}, \rho_{\overline{X}}^{\overline{U}}]$ passes through some point $\overline{y}$ lying with distance 100δ of either $\overline{u}_i$ or $\overline{u}_{i+1}$, and likewise $[\overline{y}, \rho_{\overline{X}}^{\overline{U}}]$ passes within $1200B$ of $\overline{w}_i$.

Recall that $d_X(\overline{w}_i, \overline{u}_i) \geq E_D \geq 10000B$, while $d_G(e, s_i) \leq E_M$, where E_D and E_M are independent of n . Therefore, the construction of $\overline{y}$ above is sufficient to conclude that $[\overline{y}, \overline{s}]$ and $[\overline{u}_i, \overline{y}]$ do not pass within E of $\rho_{\overline{X}}^{\overline{U}}$. From Lemma 10.6,

$$d_{\overline{U}}(\rho_{\overline{U}}^{\overline{X}}(\overline{y}), \rho_{\overline{U}}^{\overline{X}}(\overline{u}_i)) \leq O(E), \quad d_{\overline{U}}(\rho_{\overline{U}}^{\overline{X}}(\overline{u}_i), \rho_{\overline{U}}^{\overline{X}}(\overline{s})) \leq O(E).$$

We may now apply the triangle inequality and conclude that

$$d_{\overline{U}}(\rho_{\overline{U}}^{\partial(X/N)}(\chi_{i,j}), \rho_{\overline{U}}^{\overline{X}}(\overline{w}_i)) \leq O(B). \quad (4)$$

The rest of the argument follows as in Lemma 7.1. At its close, we use the new enemy prongs that were used to construct the $\overline{u}_i$ to argue as in Lemma 7.7 and obtain that, after relabeling to $\chi_1, \chi_2, \dots, \chi_Q$, we have

$$(\chi_i \cdot \chi_{i+1})_{\rho_{\overline{X}}^{\overline{U}}} \geq D' + \frac{1}{2}F.$$

Lastly, recall Equation 4. Applying the triangle equality a final time with $\overline{w}_i$ and either $\overline{x}_1$ or $\overline{x}_2$, and recalling the construction of these points, yields for all χ_i and χ_j that

$$d_{\overline{U}}(\rho_{\overline{U}}^{\partial(X/N)}(\chi_i), \rho_{\overline{U}}^{\partial(X/N)}(\chi_j)) \leq \Psi_1 + O(E).$$

We replace $O(B)$ with $O(L)$, which is suboptimal but harmless for later arguments. The proof is now complete. $\square$

Proof of Lemma 10.15. Let J be the constant in Lemma 10.14, and let $\mathcal{C}'$ be the constant in the proof of Theorem 7.8. By repeatedly applying Lemma 10.16 approximately $\frac{10(\mathcal{C}'+L)}{F}$ times, we arrive at a sequence $\xi_1 = \chi_1, \dots, \chi_Q = \xi_2$ of points in $\partial(X/N)$ with the following properties:

1. For $1 \leq i \leq Q - 1$, $(\chi_i \cdot \chi_{i+1})_{\rho_{\overline{X}}^{\overline{U}}} \geq 10C' + 10L$.
2. For all $2 \leq i \leq Q$ and $j \in \{1, 2\}$, we have

$$d_{\overline{U}}(\rho_{\overline{U}}^{\partial(X/N)}(\chi_i), \rho_{\overline{U}}^{\partial(X/N)}(\xi_j)) \leq \Psi_1 + O\left(\frac{10B(C' + L)}{F}\right).$$

We now repeat the argument of Theorem 7.8 for each pair χ_i, χ_{i+1} , substituting Lemma 10.16 for Lemma 7.1. The result is a sequence of paths $\gamma_1, \gamma_2, \dots, \gamma_{Q-1}$ in $\partial(X/N)$ with domain $[0, 1]$, each satisfying the property that for each γ_i and each $t \in [0, 1]$,

$$(\gamma(t) \cdot \chi_i)_{\rho_{\overline{X}}^{\overline{U}}} \geq (\chi_i \cdot \chi_{i+1})_{\rho_{\overline{X}}^{\overline{U}}} - C' \geq 10L$$

Applying Lemma 10.6, we conclude that

$$\text{diam}(\rho_{\overline{U}}^{\partial(X/N)}(\gamma_i)) \leq O(E)$$

for each γ_i . Concatenating these paths together and using the second property above, we may take $\Psi_2 = O\left(\frac{100B(C' + L)}{F}\right)$ and obtain the desired result. $\square$

10.2 Height, Slack, and Finiteness

From now on, assume we have a fixed open cover $\mathcal{U}$ of $\partial(X/N)$. Recall that the topology on $\partial(X/N)$ is induced by the visual metric.

Definition 10.17. Given a point $\xi \in \partial(X/N)$, we define the *height* $H(\xi)$ of ξ to be the smallest nonnegative integer so that the set

$$\{\chi \in \partial(X/N) : (\xi \cdot \chi)_{\overline{x}_0} \geq H(\xi)\}$$

is contained in an element of $\mathcal{U}$.

Given $\overline{U} \in \overline{\mathcal{D}} \setminus \{X/N\}$, we say $\overline{U}$ is *sufficiently high* if there exists a minimal integer $H(\overline{U})$ satisfying $0 \leq H(\overline{U}) \leq d_{X/N}(\overline{x}_0, \rho_{\overline{X}}^{\overline{U}}) - 100L$ so that the set

$$\{\chi \in \partial(X/N) : (\chi \cdot \rho_{\overline{X}}^{\overline{U}})_{\overline{x}_0} \geq H(\overline{U})\}$$

is contained in an element of $\mathcal{U}$. If $\overline{U}$ is sufficiently high, we define its *height* to be this quantity $H(\overline{U})$, and its *slack* $A(\overline{U})$ to be the difference $A(\overline{U}) := d_{X/N}(\overline{x}_0, \rho_{\overline{X}}^{\overline{U}}) - H(\overline{U})$.

Remark 10.18. Given any finite-diameter set $K \subseteq X$, one can define the *shadow* of K from $\overline{x}_0$ as the set of all boundary points ξ with at least one $(1, 20L)$ -quasi-geodesic representative ray $\bar{r}$ originating from $\overline{x}_0$ and passing through K . (This definition is roughly equivalent to the usual definition.) Then the slack $A(\overline{U})$ of $\overline{U}$ is coarsely the radius of the largest ball we can draw around $\rho_{\overline{X}}^{\overline{U}}$ such that the shadow of this ball is contained inside some element of $\overline{\mathcal{U}}$.

Note that by definition of the visual metric topology, all points in $\partial(X/N)$ have well-defined heights.

It will be easier to work with ‘balls’ defined in terms of Gromov products, rather than balls defined in terms of the visual metric.

Definition 10.19. Given $\xi \in \partial(X/N)$ and $D \in \mathbb{R}_{\geq 0}$, define

$$b(\xi, D) = \{\chi \in \partial X : (\xi \cdot \chi)_{\bar{x}_0} \geq D\}$$

Why is the height/slack terminology needed? Given $\xi \in \partial(X/N)$, the maximum

$$\max_{\bar{U} \in \overline{\mathcal{D}} \setminus \{X/N\}} d_{\bar{U}}(\rho_{\bar{U}}^{\bar{X}}(\bar{x}_0), \rho_{\bar{U}}^{\partial(X/N)}(\xi))$$

may be unbounded, which would block a key ‘large rotations’ step in a later argument. Slack lets us get around this problem.

Lemma 10.20. *Fix a quantity A_M , and let $\overline{\mathcal{D}}_0(A_M)$ be the set of all $\bar{U} \in \overline{\mathcal{D}}_0$ such that the slack $A(\bar{U})$ of $\bar{U}$ satisfies $A(\bar{U}) < A_M$. Then for any $\xi \in \partial X$,*

$$\max_{\bar{U} \in \overline{\mathcal{D}}_0(A_M)} d_{\bar{U}}(\bar{x}_0, \xi) < \infty.$$

Proof. By Lemma 10.6 and the Morse Lemma, the domains $\mathcal{U}$ such that some $(1, 20L)$ -quasi-geodesic ray $\bar{r}$ does not pass within distance, say, $1000L$ of $\rho_{\bar{X}}^{\bar{U}}$ may be excluded. Fixing a specific $\bar{r}$ defining a boundary point ξ , and assume $\rho_{\bar{X}}^{\bar{U}}$ does lie within $1000L$ of $\bar{r}$. Let $V \in \mathcal{U}$ be such that

$$\{\chi \in \partial(X/N) : (\xi \cdot \chi)_{\bar{x}_0} \geq H(\xi)\} \subseteq V.$$

For $\chi \in \partial(X/N)$, we have by analysis above that

$$(\chi \cdot \xi)_{\bar{x}_0} \geq \min((\chi \cdot \rho_{\bar{X}}^{\bar{U}})_{\bar{x}_0}, (\xi \cdot \rho_{\bar{X}}^{\bar{U}})_{\bar{x}_0}) - 20L$$

$$\geq \min((\chi \cdot \rho_{\bar{X}}^{\bar{U}})_{\bar{x}_0}, d_{X/N}(\bar{x}_0, \rho_{\bar{X}}^{\bar{U}})) - 2000L$$

Therefore, as long as $d_{X/N}(\rho_{\bar{X}}^{\bar{U}}, \bar{x}_0) \geq H(\xi) + 2000L$, we have the implication

$$(\chi \cdot \rho_{\bar{X}}^{\bar{U}})_{\bar{x}_0} \geq H(\xi) + 2000L \rightsquigarrow (\chi \cdot \xi)_{\bar{x}_0} \geq H(\xi).$$

By definition of $H(\xi)$, we therefore obtain

$$H(\bar{U}) \leq H(\xi) + 2000L.$$

With the additional assumption that $\bar{U} \in \bar{\mathcal{D}}_0(A_M)$, we have by definition that

$$d_{X/N}(\bar{x}_0, \rho_{\bar{X}}^{\bar{U}}) \leq A_M + H(\bar{U}) \leq A_M + H(\xi) + 2000L.$$

Furthermore, this bound holds uniformly for all $\bar{U} \in \bar{\mathcal{D}}_0(A_M)$.

Now, let $\bar{y}$ be a point on $\bar{r}$ such that $d_{X/N}(\bar{x}_0, \bar{y}) \geq A_M + H(\xi) + 10^{13}L$, so that by definition

$$d_{\bar{U}}(\rho_{\bar{U}}^{\bar{X}}(\bar{x}_0), \rho_{\bar{U}}^{\partial(X/N)}(\xi)) \leq d_{\bar{U}}(\rho_{\bar{U}}^{\bar{X}}(\bar{x}_0), \rho_{\bar{U}}^{\bar{X}}(\bar{y})) + O(n).$$

Note that we may choose $\bar{y}$ independently of $\bar{U}$. Now, by our hypotheses, there exist points $\hat{x}_0$ and $\hat{y}$ in $\mathcal{X}$ projecting respectively to $\bar{x}_0$ and $\bar{y}$. From functoriality and coarse Lipschitzness of projections, $d_{\mathcal{X}}(\hat{x}_0, \hat{y})$ provides a coarse upper bound for $d_{\bar{U}}(\rho_{\bar{U}}^{\bar{X}}(\bar{x}_0), \rho_{\bar{U}}^{\bar{X}}(\bar{y}))$.

To conclude, we note that $d_{\mathcal{X}}(\hat{x}_0, \hat{y})$ is fully independent of the $\bar{U} \in \bar{\mathcal{D}}_0(A_M)$ chosen, and so coarsely upper bounds $d_{\bar{U}}(\rho_{\bar{U}}^{\bar{X}}(\bar{x}_0), \rho_{\bar{U}}^{\bar{X}}(\bar{y}))$ for all such $\bar{U}$. This bound completes the proof. $\square$

10.3 Constructing Suitable Paths

We now move on to verifying the locally finite 1-disks property. Throughout the proof, we will use a constant $\mathcal{B}$. This constant will not depend on f_n or $\mathcal{U}$; $\mathcal{B} := 1000L$ should suffice.

The following lemma is a straightforward consequence of the triangle inequality for Gromov products.

Lemma 10.21. *Fix $\xi \in \partial(X/N)$, and let χ be either another point in $\partial(X/N)$ or a point in X/N . Assume $\chi \in b(\xi, H(\xi) + \mathcal{B})$. Then*

$$H(\chi) =_{100L} H(\xi)$$

We now prove the following lemma, which will be broken down into several subclaims. The lemma will rely on a running construction; this construction is depicted in Figures 12 and 13; note that Figure 13 is a close-up of part of 12.

Lemma 10.22. *Assume G , X , $\mathcal{X}$, and X/N satisfying Hypotheses 9.4 and 9.7. Then $\partial(X/N)$ satisfies the locally finite 1-disks property.*

Proof. We will describe the construction of the paths g_m , performing verification as we go. We begin with a path $f_m : [0, 1] \rightarrow \partial(X/N)$ and our fixed open cover $\mathcal{U}$.

The first step is a variation of the standard Lebesgue number construction. Specifically, we choose a subdivision $0 = t_0, t_1, \dots, t_k = 1$ of $[0, 1]$ and define $\xi_i = f_n(t_i)$ and $\mathcal{I}_i = f_n([t_i, t_{i+1}])$ so that for all $\mathcal{I}_i$, we have

$$\mathcal{I}_i \subseteq b(\xi_i, H(\xi_i) + 1000\mathcal{B})$$




Figure 12: A schematic of the construction used to prove Lemma 10.22, showing the relevant objects under projection to X/N , along with windows showing relevant objects in the domains $\bar{U}_{i,j}$ and $\bar{U}_{i,j+1}$. The chains of projection points of domains are such that sequential domains are pairwise orthogonal.




Figure 13: A close-up of part of Figure 10.22.

The existence of such a subdivision follows from compactness of $f_n([0, 1])$ and the definition of the $H(\cdot)$, plus an elementary topological argument. Notice the consequence that for all $1 \leq i \leq k-1$, $H(\xi_{i+1}) =_{100L} H(\xi_i)$. Since each $b(\xi_i, H(\xi_i))$ is contained in some element of $\mathcal{U}$ by construction, let V_i be this element. We also choose for each ξ_i a $(1, 20L)$ -quasi-geodesic ray $\bar{r}_i$ originating from $\bar{x}_0$ representing ξ_i . By removing unnecessary t_i , we may also assume without loss of generality that

$$(\xi_i \cdot \xi_{i+1})_{\bar{x}_0} \leq H(\xi_i) + 10000\mathcal{B}$$

We now focus on a specific ξ_i , $\mathcal{I}_i$, and V_i .

Lemma 10.23. *Given the sequences (ξ_i) , $(\mathcal{I}_i)$, and (V_i) as above, there exists for each $1 \leq i \leq k$ a sequence $\bar{U}_{i,1}, \dots, \bar{U}_{i,K_i}$ of domains in $\bar{\mathcal{D}}_0 \setminus \{X/N\}$ satisfying the following properties:*

1. The point $\rho_{\bar{X}}^{\bar{U}_{i,1}}$ lies within $2E$ of $\bar{r}_i$, and the point $\rho_{\bar{X}}^{\bar{U}_{i,K_i}}$ lies within $2E$ of $\bar{r}_{i+1}$.
2. For each $\bar{U}_{i,j}$, $\bar{U}_{i,j} \perp \bar{U}_{i,j+1}$.
3. Every $\bar{U}_{i,j}$ lies within $\mathcal{B}$ of $\bar{r}_i$.
4. For each $\bar{U}_{i,j}$, we have

$$H(\xi_i) + 100\mathcal{B} =_{10\mathcal{B}} d_{X/N}(\rho_{\bar{X}}^{\bar{U}_{i,j}}, \bar{x}_0)$$

and

$$H(\xi_i) + 100\mathcal{B} =_{10\mathcal{B}} (\xi_i \cdot \rho_{\bar{X}}^{\bar{U}_{i,j}})_{\bar{x}_0}$$

5. We have that $\bar{U}_{i,K_i} = \bar{U}_{i+1,1}$.

Proof. By construction, there exists a point $\bar{s}_i$ on $\bar{r}_i$ such that $d_{X/N}(\bar{s}_i, \bar{x}_0) =_{0.1\mathcal{B}} H(\xi_i) + 100\mathcal{B}$. We choose $\bar{s}'_i$ analogously on $\bar{r}_{i+1}$. Let $\bar{\eta} = [\bar{s}_i, \bar{s}'_i]$ be a geodesic in X/N between these two points, and let $\eta = [s_i, s'_i]$ be an arbitrary lift of $\bar{\eta}$ to X .

Since $(\xi_i \cdot \xi_{i+1})_{\bar{x}_0} \geq H(\xi_i) + 999\mathcal{B}$, $\bar{r}_i$ and $\bar{r}_{i+1}$ must fellow-travel within the range where $\bar{s}_i$ and $\bar{s}'_i$ are chosen, and we may conclude that $\bar{\eta}$, and hence η , have length at most $0.2\mathcal{B}$. In particular, every point of η lies within $\mathcal{B}$ of $\bar{r}_i$ and $\bar{r}_{i+1}$, and has distance $H(\xi_i) + 100\mathcal{B} \pm \mathcal{B}$ from $\bar{x}_0$. Applying Hypothesis 9.7, there exists a sequence $U_{i,1}, \dots, U_{i,J_i}$ of domains in $\mathcal{D}_0$ such that $\rho_X^{U_{i,1}}$ lies within E of s_i , $\rho_X^{U_{i,J_i}}$ lies within E of s_{i+1} , every $\rho_X^{U_{i+1}}$ lies within E of η , and for all $1 \leq i \leq J_i - 1$, we have $U_i \perp U_{i+1}$ and $d_X(\rho_X^{U_i}, \rho_X^{U_{i+1}}) \leq 10E$. It is now straightforward to verify that the projected sequence $\bar{U}_{i,1}, \dots, \bar{U}_{i,J_i}$ satisfies the desired properties.

Given $\bar{s}'_i$ and $\bar{s}_{i+1}$ for $1 \leq i \leq k-1$, we now build a new geodesic $\bar{\eta}$ between them. As $H(\xi_{i+1}) =_{0.1\mathcal{B}} H(\xi_i)$, this geodesic has length at most $\mathcal{B}$. We now build a sequence

$\overline{U}_{i,J_i}, \overline{U}_{i,J_i+1}, \dots, \overline{U}_{i,K_i}$ as in the previous paragraph from $\overline{\eta}$; verification of the desired coarse equalities is likewise straightforward. $\square$

Lemma 10.24. *Each $\overline{U}_{i,j}$ constructed in Lemma 10.23 has sufficient height, and its slack $A(\overline{U}_{i,j})$ is at most $200\mathcal{B}$.*

Proof. By construction, we have for each $\overline{U}_{i,j}$ that

$$H(\xi_i) + 100\mathcal{B} =_{10\mathcal{B}} d_{X/N}(\rho_{\overline{X}}^{\overline{U}_{i,j}}, \overline{x}_0)$$

and

$$H(\xi_i) + 100\mathcal{B} =_{10\mathcal{B}} (\xi_i \cdot \rho_{\overline{X}}^{\overline{U}_{i,j}})_{\overline{x}_0}$$

From Lemma 10.21, $H(\overline{U}_{i,j}) =_{0.1\mathcal{B}} H(\xi_i)$. From the control on $d_{X/N}(\rho_{\overline{X}}^{\overline{U}_{i,j}}, \overline{x}_0)$ given above, certainly $\overline{U}_{i,j}$ has sufficient height. On the other hand,

$$d_{X/N}(\overline{x}_0, \rho_{\overline{X}}^{\overline{U}_{i,j}}) - H(\overline{U}_{i,j}) \leq (H(\xi_i) + 110\mathcal{B}) - (H(\xi_i) - \mathcal{B}) \leq 200\mathcal{B}. \square$$

We now build the paths we will concatenate to construct g_n . Choose a constant $\mathcal{T}$; the exact value will be chosen in a moment. Given $\overline{U}_{i,j}$, recall that $\overline{U}_{i,j}$ contains a quasi-embedded copy of $\mathbb{R}$. As a result, there exists some point $\overline{u}_{i,j} \in \overline{U}_{i,j}$ such that $d_{\overline{U}_{i,j}}(\overline{u}_{i,j}, \rho_{\overline{U}_{i,j}}^{\overline{X}}(\overline{x}_0)) \geq 100\mathcal{T}m$. Say that the maps in Lemma 10.11 are Ψ_1 -coarsely surjective for $\Psi_1 = O(E)$. Also let J be the constant in Lemma 10.6. Using this lemma, we may construct $\lambda_{i,j} \in \partial(X/N)$ so that $\lambda_{i,j}$ $\mathcal{T}$ -coarsely projects to $(\overline{u}_{i,j}, \overline{u}_{i,j+1}) \in \overline{U}_{i,j} \times \overline{U}_{i,j+1}$. We may then apply Lemma 10.15 with Ψ_1 , letting Ψ_2 be the resulting constant. This lemma yields a path $\gamma_{i,j} : [0, 1] \rightarrow \partial(X/N)$ from $\lambda_{i,j}$ to $\lambda_{i,j+1}$ such that for all $t \in [0, 1]$,

$$d_{\overline{U}_{i,j}}(\rho_{\overline{U}_{i,j}}^{\partial(X/N)}(\gamma_{i,j}(t)), \rho_{\overline{U}_{i,j}}^{\partial(X/N)}(\lambda_{i,j})) \leq \Psi_2$$

We obtain finitely many such $\gamma_{i,j}$, and so may concatenate them together into a single path $g_n : [0, 1] \rightarrow \partial(X/N)$ so that $g_n|_{[t_i, t_{i+1}]}$ is equal to the concatenation of $\gamma_{i,1}$ through γ_{i,K_i-1} .

We now require that $\mathcal{T} \geq 100(\Psi_1 + \Psi_2 + J)$.

It remains to verify that g_n satisfies the two properties of Definition 9.2. Recall that for $t \in [t_i, t_{i+1}]$, by definition $f_m(t) \in \mathcal{I}_i \subseteq V_i$.

Lemma 10.25. *For each $\gamma_{i,j}$, $\gamma_{i,j}([0, 1]) \subseteq V_i$.*

Proof. For each $t \in [0, 1]$, we have

$$d_{\overline{U}_{i,j}}(\rho_{\overline{U}_{i,j}}^{\partial(X/N)}(\gamma_{i,j}(t)), \rho_{\overline{U}_{i,j}}^{\partial(X/N)}(\lambda_{i,j})) \leq \Psi_2$$

as well as

$$d_{\overline{U}_{i,j}}(\rho_{\overline{U}_{i,j}}^{\partial(X/N)}(\lambda_{i,j}), \rho_{\overline{U}_{i,j}}^{\partial(X/N)}(\overline{u}_{i,j})) \leq \Psi_1.$$

But also

$$d_{\overline{U}_{i,j}}(\overline{u}_{i,j}, \rho_{\overline{U}_{i,j}}^{\overline{X}}(\overline{x}_0)) \geq 100\mathcal{T}m.$$

By choice of $\mathcal{T}$ and the triangle inequality, we obtain

$$d_{\overline{U}_{i,j}}(\rho_{\overline{U}_{i,j}}^{\partial(X/N)}(\gamma_{i,j}(t)), \rho_{\overline{U}_{i,j}}^{\overline{X}}(\overline{x}_0)) \geq 10J.$$

Therefore, Lemma 10.3 implies that any $(1, 20L)$ -quasi-geodesic ray originating at $\overline{x}_0$ representing $\gamma_{i,j}(t)$ passes within J of $\rho_{\overline{X}}^{\overline{U}_{i,j}}$. The fact that $(\xi_i \cdot \rho_{\overline{X}}^{\overline{U}_{i,j}})_{\overline{x}_0} =_{10\mathcal{B}} H(\xi_i) + 100\mathcal{B}$ implies that, for $\mathcal{B} \geq 10J + 1000L$, we have

$$(\xi_i \cdot \gamma_{i,j}(t))_{\overline{x}_0} \geq H(\xi_i) + 50\mathcal{B}$$

By definition of $H(\xi_i)$, this inequality is sufficient to conclude that $\gamma_{i,j}(t) \in V_i$. $\square$

Note for the following lemma that any $b(\chi, R)$ with $\chi \in \partial(X/N)$ and $R \in \mathbb{R}_{>0}$ contains some open neighborhood of ξ , essentially by definition of the visual metric topology.

Lemma 10.26. *For each point $\xi \in \partial(X/N)$, there is some R_0 and M_0 so that for $m \geq M_0$, g_m does not intersect $b(\xi, R_0)$, where*

$$b(\xi, R_0) = \{\chi \in \partial(X/N) : (\xi \cdot \chi)_{\overline{x}_0} \geq R_0\}$$

Proof. Define

$$\mu := \max_{\overline{U} \in \overline{\mathcal{D}}_0(10^5\mathcal{B})} d_{\overline{U}}(\rho_{\overline{X}}^{\partial(X/N)}(\xi), \rho_{\overline{U}}^{\overline{X}}(\overline{x}_0))$$

By Lemma 10.20, μ is finite. Choose $M_0 \geq \mu + 10E$, and $R_0 = H(\xi) + 10^{10}\mathcal{B}$.

For every $t \in [0, 1]$ and $m \geq M_0$, by construction there exists some $\overline{U} \in \overline{\mathcal{D}}_0(200\mathcal{B})$ (one of the $\overline{U}_{i,j}$ for that g_m , with slack at most $200\mathcal{B}$) such that

$$d_{\overline{U}}(\rho_{\overline{U}}^{\overline{X}}(\overline{x}_0), \rho_{\overline{X}}^{\partial(X/N)}(g_m(t))) \geq 100\mathcal{T}M_0.$$

On the other hand, by definition of M_0 ,

$$d_{\overline{U}}(\rho_{\overline{U}}^{\overline{X}}(\overline{x}_0), \rho_{\overline{U}}^{\partial(X/N)}(\xi)) \leq M_0.$$

From the triangle inequality, we may conclude that

$$d_{\overline{U}}(\rho_{\overline{U}}^{\partial(X/N)}(\xi), \rho_{\overline{X}}^{\partial(X/N)}(g_m(t))) \geq \mathcal{T}$$

. As $\mathcal{T}$ was chosen sufficiently large, we may apply Lemma 10.6 and obtain

$$(g_m(t) \cdot \xi)_{\rho_{X/N}^{\bar{U}}} \leq E.$$

If $g_m(t) \in b(\xi, R_0)$, we would have

$$(g_m(t) \cdot \xi)_{\bar{x}_0} \geq H(\xi) + 10^{10}\mathcal{B}$$

by definition. By standard properties of the Gromov product, we must therefore have that

$$d_{X/N}(\bar{x}_0, \rho_{\bar{X}}^{\bar{U}}) \geq H(\xi) + 10^9\mathcal{B}$$

and that

$$(\rho_{\bar{X}}^{\bar{U}} \cdot \xi)_{\bar{x}_0} \geq H(\xi) + 10^9\mathcal{B}.$$

From Lemma 10.21, we have that $H(\bar{U}) \leq H(\xi) + \mathcal{B}$, and so by definition the slack $A(\bar{U})$ of $\bar{U}$ is at least $10^8\mathcal{B}$. This slack bound contradicts the assumption that $\bar{U} \in \bar{\mathcal{D}}_0(200\mathcal{B})$, so in fact $g_m(t)$ cannot intersect $b(\xi, R_0)$. $\square$

We conclude that the g_m we have constructed satisfy the desired conditions, and so $\partial(X/N)$ has the locally finite 1-disks property, thus proving Lemma 10.22. $\square$

11 The Nöbeling Curve and Applications

Assembling the results of the last few sections, we have the following Theorem.

Theorem 11.1. *Let G , $\mathcal{X}$, X , and N be as in Hypothesis 9.7. Additionally assume that $\dim(\partial X) = 0$. Then $\partial(X/N)$ is homeomorphic to the Nöbeling curve.*

Note that all quasi-trees have zero-dimensional boundaries.

We now identify two examples of classes of groups and spaces satisfying these conditions.

11.1 Graphs of Groups

The first class comes from [BR22], although the reader should note that this work builds on ideas developed around graphs of groups and around cube complexes by many researchers (see, e.g., [BHS17]).

Definition 11.2. Let Γ be a finite simplicial graph with vertex set V_Γ and edge set E_Γ . For each $v \in V_\Gamma$, let G_v be a nontrivial finitely generated group. We define the *graph product* G_Γ as the group

$$G_\Gamma := \left(\bigast_{v \in V_\Gamma} G_v \right) / \langle\langle [g_v, g_w] : g_v \in G_v, g_w \in G_w, \{v, w\} \in E_\Gamma \rangle\rangle.$$

Given an induced subgraph $\Lambda \subseteq \Gamma$, we have that G_Λ is a subgroup of G_Γ ; we call such groups *graph subgroups*.

With this notation, right-angled Artin groups (RAAGs) are exactly the graph products G_Γ where $G_v = \mathbb{Z}$ for all vertices v .

Berlyne and Russell prove that graph products are relative HHGs, with explicitly constructed HHS spaces realizing the HHG actions.

Theorem 11.3 (Berlyne-Russell [BR22], Theorem A). *Let G_Γ be a graph product of groups. Then G_Γ is relatively hierarchically hyperbolic relative to its vertex groups, with explicit maximal domain X_G . In particular, if all vertex groups of G_Γ are hyperbolic or hierarchically hyperbolic, then G_Γ is hierarchically hyperbolic.*

We will now identify the class of graph products to which the Nöbeling curve result will apply.

Definition 11.4. Let Γ be a finite simplicial graph, and let $\Lambda_1, \dots, \Lambda_k$ be nonempty induced subgraphs satisfying $\Lambda_i \neq \Gamma$. Set $\hat{\Lambda} = \{\Lambda_1, \dots, \Lambda_k\}$. Recall that the link $\text{lk}(\Lambda)$ of an induced subgraph $\Lambda \subseteq \Gamma$ is the set of all vertices of Γ that are adjacent or equal to every vertex of Λ . The *collapse graph* $\Gamma_{\hat{\Lambda}}$ of Γ relative to $\hat{\Lambda}$ is the graph with vertex set Λ and Λ_i and Λ_j connected by an edge if $\Lambda_i \subseteq \text{lk}(\Lambda_j)$.

Definition 11.5. Let G_Γ be a graph product of groups. We say that G_Γ is *admissible* if there exists a collection $\hat{\Lambda} = \{\Lambda_1, \dots, \Lambda_k\}$ of induced proper subgraphs such that the following conditions are satisfied:

1. Γ is not a nontrivial join.
2. Each $\Lambda_i \in \hat{\Lambda}$ is either a single vertex v such that G_v possesses some finitely generated Cayley graph containing a quasi-isometrically embedded copy of $\mathbb{R}$, or is a subgraph with at least two vertices that is not a join.
3. Every vertex of Γ lies in some Λ_i .
4. The collapse graph $\Gamma_{\hat{\Lambda}}$ is connected.

Remark 11.6. Note that in Definition 11.5, if all G_v satisfy the embedded- $\mathbb{R}$ condition, then it suffices for G_Γ to be connected and to not be a join, as $\hat{\Lambda}$ can be set to be the vertex set of Γ .

Lemma 11.7. *Let G_Γ be a graph product of groups with vertex set V_Γ . Let N be as in hypothesis 3.1, and let X_G be the top-level domain obtained in Theorem 11.3. Assume that G_Γ is admissible.*

Then $\partial(X_G/N)$ is homeomorphic to the Nöbeling curve.

To prove this lemma, we need to prove that the setup satisfies Hypotheses 9.7 and 9.8, which require digging into how the relative HHG structure for G_Γ is constructed. For each $v \in \Gamma$, let S_v be a finite generating set of G_v , and define $S = \bigcup_{v \in \Gamma} S_v$.

Definition 11.8 ([BR22], Defs. 3.1 and 3.3.). Let G_Γ be a graph product. The domains of G_Γ are indexed by cosets of the form gG_Λ , where $g \in G_\Gamma$ and G_Λ is a graph subgroup, up to the equivalence relation that $[gG_\Lambda] = [g'G_{\Lambda'}]$ if $\Lambda = \Lambda'$ and $g^{-1}g' \in \langle G_v : v \in \Lambda \cup \text{lk}(\Lambda) \rangle$. Inclusion is defined such that Γ is the unique maximal domain.

For each induced subgraph $\Lambda \subseteq \Gamma$, define its domain $\mathcal{C}(\Lambda)$ to be the graph with vertex set G_Λ and edges between vertices $g, h \in G_\Lambda$ exactly when either $g^{-1}h \in \bigcup_{v \in \Lambda} S_v$ or there exists a proper induced subgraph $\Lambda' \subsetneq \Lambda$ such that $g^{-1}h \in G_{\Lambda'}$. Then define $\mathcal{C}([g\Lambda]) = g\mathcal{C}(\Lambda)$.

Definition 11.9 ([BR22], Def. 3.17). Given a non-maximal domain $g\Lambda$, we define

$$\rho_\Gamma^{g\Lambda} = \bigcup_{[h\Lambda] = [g\Lambda]} hG_\Lambda$$

We will often somewhat cheat on notation and use $[g\Lambda]$ as a standin for $\rho_\Gamma^{[g\Lambda]}$ when it is clear from context that we are discussing the image $\rho_\Gamma^{[g\Lambda]}$.

Definition 11.10 ([BR22], Def. 3.19). Given two domains $[g_1\Lambda_1], [g_2\Lambda_2]$, we define $[g_1\Lambda_1] \perp [g_2\Lambda_2]$ if there exists $h \in G_\Gamma$ with $[h\Lambda_1] = [g_1\Lambda_1]$, $[h\Lambda_2] = [g_2\Lambda_2]$, and $\Lambda_2 \in \text{lk}(\Lambda_1)$.

We now prove Lemma 11.7.

Proof of Lemma 11.7. We need to establish Hypotheses 9.7 and 9.8. Recall that the top-level domain for G_Γ is $\mathcal{C}(\Gamma)$. From Theorems 4.4 and 4.10 in [BR22], $\mathcal{C}(\Gamma)$ is an infinite-diameter quasi-tree, establishing Hypothesis 9.8.

We now turn to 9.7. From 11.3 and standard facts about relative HHGs, it suffices to verify the large dense domains property. Let $\hat{\Lambda} = \{\Lambda_1, \dots, \Lambda_k\}$ be a set verifying the assumption that G_Γ is admissible. We define

$$\mathcal{D}_0 = \{[g\Lambda_i] : g \in G_\Gamma, \Lambda_i \in \hat{\Lambda}\}.$$

As the G_Λ -action on domains satisfies $g[h\Lambda] = [gh\Lambda]$, certainly $\mathcal{D}_0$ is closed under this action. Furthermore, for every $\mathcal{C}([g\Lambda_i])$, either $\Lambda_i = \{v\}$ and $\mathcal{C}([g\Lambda_i])$ is isometric to the Cayley graph $\text{Cay}(G_v, S_v)$, or Λ_i is not a join. In the first case, $\text{Cay}(G_v, S_v)$ contains a quasi-isometrically embedded copy of $\mathbb{R}$ by assumption, and in the second case, we may apply Theorems 4.4 and 4.10 in [BR22] to conclude that $\mathcal{C}([g\Lambda_i])$ is an infinite-diameter quasi-tree, which necessarily contains a quasi-isometrically embedded copy of $\mathbb{R}$.

Next, for any $g \in \mathcal{C}(\Gamma)$, g lies within gG_v for any $v \in \Gamma$, and hence g lies within $[g\Lambda_i]$ for any $[g\Lambda_i]$.

Lastly, it suffices to show that there is some constant $\mathcal{R}$ with the following property: for any adjacent vertices $g, h \in \mathcal{C}(\Gamma)$ and any $[g_1\Lambda_{j_1}]$ containing g , there exists a sequence $[g_1\Lambda_{j_1}], \dots, [g_q\Lambda_{j_q}]$ so that $[g_i\Lambda_{j_i}] \perp [g_{i+1}\Lambda_{j_{i+1}}]$ for all $[g_i\Lambda_{j_i}]$, all $[g_i\Lambda_{j_i}]$ have distance at most $\mathcal{R}$ from g , and $[g_q\Lambda_{j_q}]$ contains h .

One possibility is that $g^{-1}h = s_v$ for some $s_v \in S_v$, where S_v is the generating set of some vertex group G_v . Assume without loss of generality that $v \in \Lambda_1$. Then $h \in [g\Lambda_1]$, as in fact $h \in gG_{\Lambda_1}$. On the other hand, $g \in [g_1\Lambda_{j_1}]$ by assumption, so by definition there exists g'_1 such that $[g'_1\Lambda_{j_1}] = [g_1\Lambda_{j_1}]$ and $g \in g'_1G_{\Lambda_{j_1}}$. So in fact $[g_1\Lambda_{j_1}] = [g\Lambda_{j_1}]$.

Now let γ be a path in the collapse graph $\Gamma_{\hat{\Lambda}}$ from Λ_{j_1} to Λ_1 ; here we need the assumption that the collapse graph $\Gamma_{\hat{\Lambda}}$ is connected. Enumerate γ as $\gamma = \{\Lambda_{j_1}, \Lambda_{j_2}, \dots, \Lambda_{j_q}\}$ with $\Lambda_{j_q} = \Lambda_1$. Considering the sequence $[g\Lambda_{j_1}, g\Lambda_{j_2}, \dots, g\Lambda_{j_q}]$, by definition of the collapse graph, we have $[g\Lambda_{j_i}] \perp [g\Lambda_{j_{i+1}}]$ for all $[g\Lambda_{j_i}]$. Furthermore, all these domains clearly have distance 0 from g , while the last domain has distance 0 from h by choice of $\Lambda_{j_q} = \Lambda_1$.

The other case we must consider is when $g^{-1}h \in G_{\Lambda'}$ for some $\Lambda' \subsetneq \Gamma$. Say $g^{-1}h = s$, and write s as a product of generators $s_1s_2\dots s_{p'}$, where each $s_i \in G_{\Lambda'}$. Then $gs_1s_2\dots s_{p'}$ is adjacent to h for all $1 \leq p' \leq p$, and we may rerun the previous argument for each s_i and then concatenate the resulting sequences of domains. $\square$

11.2 Non-Geometric Three-Manifold Groups

The second class comes from [Hag+25], which concerns 3-manifold groups. Our starting point is the following theorem:

Theorem 11.11 (Hagen-Russell-Sisto-Spriano, Corollary 1.3). *If M is a 3-dimensional non-geometric graph manifold, then $\pi_1(M)$ is a hierarchically hyperbolic group, with explicit top-level domain X_M .*

We have the following result.

Lemma 11.12. *Let M be a 3-dimensional non-geometric graph manifold, and let X_M be the top-level domain obtained in Theorem 11.11. Let N be as in Hypothesis 3.1.*

Then $\partial(X_M/N)$ is homeomorphic to the Nöbeling Curve.

As before, we need to check Hypotheses 9.7 and 9.8. We note the following facts.

Lemma 11.13. *Let M be a 3-dimensional non-geometric graph manifold, and let X_M be the top-level domain obtained in Theorem 11.11. Then the following properties hold.*

1. *There exists a natural quasi-isometry $\sigma : T \rightarrow X_M$, where T is an infinite-diameter tree (the Bass-Serre tree for $\pi_1(M)$).*
2. *There exists a subcollection of non-maximal domains corresponding to vertices v of T , and for each such vertex v , the corresponding hyperbolic space is a quasi-line L_v .*
3. *Furthermore, the projection of the domain corresponding to $v \in T$ to X_M contains the point $\sigma(v)$.*
4. *For vertices $v, s \in T$, the domains corresponding to v and s are orthogonal if and only if v and s are adjacent in T .*

Proof. The first three results follow from Lemma 6.8 in [Hag+25], although the reader should note that the authors use a combinatorial definition of hierarchical hyperbolicity derived from the combinatorial HHG machinery developed in [Beh+24], and some care must be taken to unravel the relevant definitions. The fourth result follows from the definition of orthogonality in [Beh+24] and Lemma 6.2 in [Hag+25]. $\square$

Lemma 11.12 now follows directly from Lemma 11.13.

References

- [Abb+26] Carolyn Abbott, Daniel Berlyne, Giorgio Mangioni, Thomas Ng, and Alexander J. Rasmussen. *Random quotients preserve acylindrical and hierarchical hyperbolicity*. 2026. arXiv: 2507.16677 [math.GR]. URL: <https://arxiv.org/abs/2507.16677>.
- [AW26] Daniel Allcock and Alex Wright. *Hierarchical hyperbolicity, ball quotients, and the moduli space of smooth cubic surfaces*. 2026. arXiv: 2609.14416 [math.GR]. URL: <https://arxiv.org/abs/2609.14416>.
- [BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. “Constructing group actions on quasi-trees and applications to mapping class groups”. In: *Publ. Math., Inst. Hautes Étud. Sci.* 122 (2015), pp. 1–64.
- [BCM12] Jeffrey Brock, Richard Canary, and Yair Minsky. “The classification of Kleinian surface groups II: The ending lamination conjecture”. In: *Ann. of Math.* 176.2 (2012), pp. 1–149.
- [Beh+12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. “Geometry and rigidity of mapping class groups”. In: *Geom. Topol.* 16.2 (2012), pp. 781–888.
- [Beh+24] Jason Behrstock, Mark F. Hagen, Alexandre Martin, and Alessandro Sisto. “A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups”. In: *Topology* 17.3 (2024).
- [Bes+19] Mladen Bestvina, Ken Bromberg, Koji Fujiwara, and Alessandro Sisto. “Acylindrical actions on projection complexes”. In: *Enseign. Math.* 65.1-2 (2019), pp. 1–32.
- [BF02] Mladen Bestvina and Koji Fujiwara. “Bounded cohomology of subgroups of mapping class groups”. In: *Geom. Topol.* 6 (2002), pp. 69–89.
- [BHS17] Jason Behrstock, Mark Hagen, and Alessandro Sisto. “Hierarchically hyperbolic spaces I: Curve complexes for cubical groups”. In: *Geom. Topol.* 21.3 (2017), pp. 1731–1804.

- [BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto. “Hierarchically hyperbolic spaces II: Combination theorems and the distance formula”. In: *Pacific J. Math.* 299 (2019), pp. 257–338.
- [BM08] Jason Behrstock and Yair Minsky. “Dimension and rank for mapping class groups”. In: *Ann. of Math.* 167.2 (2008), pp. 1055–1077.
- [BM91] Mladen Bestvina and Geoffrey Mess. “The boundary of negatively curved groups”. In: *J. Amer. Math. Soc.* 4.3 (1991), pp. 469–481.
- [Bou+23] Adrien Boulanger, Pierre Mathieu, Cagri Sert, and Alessandro Sisto. “Large deviations for random walks on hyperbolic spaces”. In: *Ann. Sci. de l’École Normale Sup.* (2023).
- [Bow98] Brian Bowditch. “Cut points and canonical splittings of hyperbolic groups”. In: *Acta Math.* 180.2 (1998), pp. 145–186.
- [BR22] Daniel Berlyne and Jacob Russell. “Hierarchical hyperbolicity of graph products”. In: *Groups Geom. Dyn.* 16 (2022), pp. 523–580.
- [CH17] Matthew Cordes and David Hume. “Stability and the Morse boundary”. In: *Jour. London Math. Soc.* 95.3 (2017), pp. 963–988.
- [Cha95] Christophe Champetier. “Statistical Properties of Groups with Finite Presentations”. In: *Adv. Math.* 116.2 (1995), pp. 197–262.
- [CM22] Matt Clay and Johanna Mangahas. “Hyperbolic quotients of projection complexes”. In: *Groups Geom. Dyn.* 16.1 (2022), pp. 225–246.
- [CMM21] Matt Clay, Johanna Mangahas, and Dan Margalit. “Right-angled Artin groups as normal subgroups of mapping class groups”. In: *Compositio Math.* 157.8 (2021), pp. 1807–1852.
- [Cou16] Rémi Coulon. “Partial periodic quotients of groups acting on a hyperbolic space”. In: *Ann. Inst. Fourier* 66.6 (2016), pp. 1773–1857.
- [CT07] James Cannon and William Thurston. “Group invariant Peano curves”. In: *Geom. Topol.* 11.3 (2007), pp. 1315–1356.
- [DGO16] François Dahmani, Vincent Guirardel, and Denis Osin. “Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces”. In: *Mem. Amer. Math. Soc.* 245.1156 (2016).
- [DGP11] François Dahmani, Vincent Guirardel, and Piotr Przytycki. “Random groups do not split”. In: *Math. Ann.* 349 (2011), pp. 657–673.
- [DMS25] Matthew Durham, Yair Minsky, and Alessandro Sisto. *Asymptotically CAT(0) metrics, Z-structures, and the Farrell-Jones Conjecture*. 2025. arXiv: 2504.17048 [math.GT]. URL: <https://arxiv.org/abs/2504.17048>.

- [Gab09] David Gabai. “Almost filling laminations and the connectivity of ending lamination space”. In: *Geom. Topol.* 13.2 (2009), pp. 1017–1041.
- [Gou22a] Sébastien Gouëzel. “Exponential bounds for random walks on hyperbolic spaces without moment conditions”. In: *Tunis. J. Math.* 4.4 (2022), pp. 635–671. ISSN: 2576-7658. DOI: 10.2140/tunis.2022.4.635. URL: <https://doi.org/10.2140/tunis.2022.4.635>.
- [Gou22b] Sébastien Gouëzel. “Exponential bounds for random walks on hyperbolic spaces without moment conditions”. In: *Tunisian Jour. Math.* 4.4 (2022), pp. 635–671.
- [Gro93] Mikhail Gromov. “Asymptotic invariants of infinite groups”. In: *Geometric Group Theory, Vol. 2*. Vol. 182. Cambridge Univ. Press, 1993, pp. 1–295.
- [Hag+25] Mark Hagen, Jacob Russell, Alessandro Sisto, and Davide Spriano. “Equivariant Hierarchically Hyperbolic Structures for 3-Manifold Groups via Quasimorphisms”. In: *Ann. Inst. Fourier* 75.2 (2025), pp. 769–828.
- [HMS22] Mark Hagen, Alexandre Martin, and Alessandro Sisto. “Extra-large type Artin groups are hierarchically hyperbolic”. In: *Math. Ann.* 388 (2022), pp. 867–938.
- [HP11] Sebastian Hensel and Piotr Przytycki. “The ending lamination space of the five-punctured sphere is the Nöbeling curve”. In: *J. London Math. Soc.* 84.2 (2011), pp. 103–119.
- [Ioa26] Rafaela Ioannou. *Menger Curve and Non-Planar Boundaries in Quotients of Hyperbolic Groups*. 2026. arXiv: 2610.03411 [math.GR]. URL: <https://arxiv.org/abs/2610.03411>.
- [KK11] Michael Kapovich and Bruce Kleiner. “Hyperbolic groups with low-dimensional boundary”. In: *Ann. Sci. École Norm. Sup. (4)* 33.5 (2011), pp. 647–669.
- [Kla22] Erica Klarreich. “The boundary at infinity of the curve complex and the relative Teichmüller space”. In: *Groups Geom. Dyn.* 16 (2022), pp. 705–723.
- [Kle07] Bruce Kleiner. “The asymptotic geometry of negatively curved spaces: uniformization, geometrization and rigidity”. In: *Proceedings of the International Congress of Mathematicians Madrid*. Vol. Volume II: Invited Lectures. Eur. Math. Soc. Press, 2007.
- [KLT97] Kazuhiro Kawamura, Michael Levin, and Edward Tymchatyn. “A characterization of 1-dimensional Nöbeling spaces”. In: *Proceedings of the 12th Summer Conference on General Topology and its Applications (North Bay, ON)*. 1997, pp. 155–174.
- [LMS11] Christopher Leininger, Mahan Mj, and Saul Schleimer. “The universal Cannon-Thurston map and the boundary of the curve complex”. In: *Comment. Math. Helv.* 86.4 (2011), pp. 769–816.

- [LS09] Christopher Leininger and Saul Schleimer. “Connectivity of the space of ending laminations”. In: *Duke Math J.* 150.3 (2009), pp. 533–575.
- [Min06] Yair Minsky. “Curve complexes, surfaces and 3-manifolds”. In: *Proceedings of the International Congress of Mathematicians*. Vol. 2. Zürich: Eur. Math. Soc., 2006, pp. 1001–1033.
- [MM00] Howard Masur and Yair Minsky. “Geometry of the complex of curves. II. Hierarchical Structure”. In: *Geom. Func. Anal.* 10.4 (2000), pp. 902–974.
- [MM99] Howard Masur and Yair Minsky. “Geometry of the complex of curves. I. Hyperbolicity”. In: *Invent. Math.* 138.1 (1999), pp. 103–149.
- [MS19] Joseph Maher and Alessandro Sisto. “Random subgroups of acylindrically hyperbolic groups and hyperbolic embeddings”. In: *Int. Math. Res. Not.* 2019.13 (2019), pp. 3941–3980.
- [MT18] Joseph Maher and Giulio Tiozzo. “Random walks on weakly hyperbolic groups”. In: *J. Reine Angew. Math.* 742 (2018), pp. 187–239.
- [MT21] Joseph Maher and Giulio Tiozzo. “Random walks, WPD actions, and the Cremona group”. In: *Proc. Lond. Math. Soc.* 123.2 (2021), pp. 153–202.
- [Osi16] Denis Osin. “Acylindrically hyperbolic groups”. In: *Trans. Amer. Math. Soc.* 368.2 (2016), pp. 851–888.
- [Pap95] P. Papasoglu. “Strongly geodesically automatic groups are hyperbolic”. In: *Invent. Math.* 121.2 (1995), pp. 323–334. ISSN: 0020-9910. DOI: 10.1007/BF01884301. URL: <https://doi.org/10.1007/BF01884301>.
- [Raf07] Kasra Rafi. “A combinatorial model for the Teichmüller metric”. In: *Geom. Funct. Anal.* 17.3 (2007), pp. 936–959.
- [Wri24] Alex Wright. *Spheres in the curve graph and linear connectivity of the Gromov boundary*. 2024. arXiv: 2304.03004 [math.GT]. URL: <https://arxiv.org/abs/2304.03004>.