Random quotients preserve acylindrical and hierarchical hyperbolicity

Carolyn Abbott, Daniel Berlyne, Giorgio Mangioni,
Thomas Ng, and Alexander J. Rasmussen

Abstract

We show that random quotients of acylindrically hyperbolic groups, obtained by taking a quotient of the group by the nth steps of a finite collection of independent random walks, are again acylindrically hyperbolic asymptotically almost surely. Our main tools come from spinning families and projection complexes, which we relate to random walks and develop further. Furthermore, we show that a random quotient of a hierarchically hyperbolic group is again hierarchically hyperbolic asymptotically almost surely. The same techniques also yield that a random quotient of a non-elementary hyperbolic group (relative to a finite collection of peripheral subgroups) is asymptotically almost surely hyperbolic (relative to isomorphic peripheral subgroups).

Oh, many a shaft at random sent
Finds mark the archer little meant!
And many a word at random spoken
May soothe, or wound, a heart that's
broken!

Sir Walter Scott, *Lord of the Isles*

Contents

1 Introduction

Two important classes of groups with strong negative curvature features are acylindrically hyperbolic groups and hierarchically hyperbolic groups. Acylindrically hyperbolic groups were introduced by Osin [Osi16] as a generalization of non-elementary hyperbolic and relatively hyperbolic groups. This broad class includes non-exceptional mapping class groups, non-virtually cyclic CAT(0) groups that do not split as direct products, $\mathrm{Out}(\mathbb{F}_n)$ for $n \geqslant 2$, the Cremona group, and any group that admits a presentation with at least two more generators than relators. Hierarchically hyperbolic groups (HHG) were introduced by Behrstock, Hagen, and Sisto in [BHS17b] to generalize the subsurface projection machinery of mapping class groups [MM99, MM00] to a wider class of groups, including right-angled Artin and Coxeter groups [BHS17b] and, more generally, most cubulated groups, and most 3-manifold groups [BHS19, HRSS25]. There is a large overlap between these two classes of groups: any HHG that is not quasi-isometric to a non-trivial product or a line is acylindrically hyperbolic. While acylindrical hyperbolicity is known to be fairly common, evidence has recently been mounting that hierarchical hyperbolicity is also widespread, as it has been shown for a large class of Artin groups, graph products, lattices, and group extensions [BHS17a, BR22, HMS24, Hug22, HV24, Rus21, FFMS25].

In this direction, one can ask which quotients of an acylindrically hyperbolic group (resp., HHG) are themselves acylindrically hyperbolic (resp., HHG). Indeed, quotients are a common tool for constructing negatively curved groups. Random groups in both the few relators model and the density model (with small enough density) are hyperbolic with overwhelming probability [Gro93, Ol'92, Oll04]. Moreover, Delzant showed that quotients of hyperbolic groups by elements with large translation length are again hyperbolic [Del96]. In fact, Groves and Manning, and independently Osin, generalized Thurston's hyperbolic Dehn filling theorem to show that all peripheral quotients of relatively hyperbolic groups whose kernels avoid a finite set of elements are again relatively hyperbolic. These quotients are also typically hyperbolic; see [GM08, Theorem 7.2] and [Osi07, Theorem 1.1].

This paper is a continuations of the above themes. More precisely, we consider quotients by the nth steps of finitely many independent random walks $w_{1,n}, \ldots, w_{k,n}$ associated to *permissible* probability measures $\mu_1, \ldots, \mu_k$. We postpone the definition of a permissible probability measure to Definition 5.1, but the reader should have in mind the case when μ is supported on a finite, symmetric generating set for G. The quotient $G/\langle\langle w_{1,n}, \ldots, w_{k,n}\rangle\rangle$ is a *random quotient* of G. Given a property P, we say that a random quotient of G has property P *asymptotically almost surely* (a.a.s.) if the probability that $G/\langle\langle w_{1,n}, \ldots, w_{k,n}\rangle\rangle$ has property P approaches 1 as n tends to infinity.

Theorem A. *Let G be an acylindrically hyperbolic group, and let $\mu_1, \dots, \mu_k$ be permissible probability measures on G. A random quotient of G is a.a.s. acylindrically hyperbolic.*

Moreover, when G is hierarchically hyperbolic, we obtain a stronger structural result for a random quotient, providing further evidence that hierarchically hyperbolic groups are common:

Theorem B. *Let G be an acylindrically hyperbolic (relative) HHG, and consider k permissible probability measures $\mu_1, \dots, \mu_k$ on G. A random quotient of G is a.a.s. an acylindrically hyperbolic (relative) HHG.*

Considering only acylindrically hyperbolic HHGs is necessary, as product HHGs include groups such as the Burger–Mozes group, which is simple and so has no non-trivial quotients.

Theorem B is already new for random quotients of mapping class groups. In this setting, the explicit description of the HHG structure has already been used by the third author to prove that random quotients of mapping class groups are quasi-isometrically rigid, in the sense that if a finitely generated group is quasi-isometric to such a quotient, then it is weakly commensurable to it [Man23]. Other hierarchically hyperbolic quotients of mapping class groups include quotients by suitable powers of all Dehn twists [BHMS24]; by suitable powers of a pseudo-Anosov element, for surfaces without boundary [BHS17a] and for the four-strand braid group [FFMS25]; and by deep enough subgroups of certain convex-cocompact subgroups [BHS17a].

In the case of non-elementary hyperbolic groups, we can deduce stronger results about the random quotient.

Corollary C. *Let G be a non-elementary hyperbolic group, and let $\mu_1, \dots, \mu_k$ be permissible probability measures on G. A random quotient of G is a.a.s. non-elementary hyperbolic.*

To the best of our knowledge, this result is not explicitly written down in the literature. However, it follows quickly from a theorem of Delzant [Del96, Théorème I] (originally stated by Gromov [Gro87, Theorem 5.5.D]), combined with the fact that the translation length of each $w_{i,n}$ is a.a.s. linear in n.

Using that Theorem B holds for the larger class of *relative* HHGs, which includes relatively hyperbolic groups, we also show that random quotients preserve non-elementary relative hyperbolicity:

Corollary D. *Let G be a group that is non-elementary hyperbolic relative to a collection $\{H_1, \dots H_\ell\}$ of finitely generated, infinite subgroups. If $\overline{G}$ is a random quotient of G, then the following hold a.a.s.*

1. *Each H_i embeds in $\overline{G}$, with image $\overline{H}_i$.*
2. *The quotient $\overline{G}$ is non-elementary hyperbolic relative to $\{\overline{H}_1, \dots, \overline{H}_\ell\}$.*

It was recently shown that a version of Corollary D holds for random quotients of free products using a density model of randomness [EMM$^+$25, Theorem 1.2]. Although the notions of randomness are different, both frameworks for random quotients yield relatively hyperbolic structures on the quotient with the same peripheral subgroups. In contrast, Dehn filling quotients of relatively hyperbolic groups are obtained by taking the quotient by sufficiently deep subgroups of the peripherals, and therefore never preserve the peripheral structure.

In [Man24], the third author introduced *short HHGs*, a particularly simple class of non-relatively-hyperbolic HHGs, which includes extra-large type Artin groups, numerous RAAGs, and non-geometric graph manifolds groups. In the spirit of Corollary C and Corollary D, we expect that one could deduce from Theorem B that random quotients of short HHGs are themselves short HHGs a.a.s. As a consequence, such quotients would be fully residually hyperbolic [MS24, Corollary H].

Outline of the arguments.

For much of the paper, we work in the following general context. Let G be a group with a non-elementary action on a hyperbolic space X. Assume there exist quasiconvex, geometrically separated subspaces $Y_1, \dots, Y_k \subseteq X$ with metrically proper, cobounded actions of subgroups $H_1, \dots, H_k$, and denote by $\widehat{X}$ the *cone-off* of X with respect to the translates of the Y_i. These assumptions allow one to build a *projection complex*, as defined by Bestvina, Bromberg, and Fujiwara [BBF15], with respect to the translates of Y_i; see Section 2.3 for generalities on projection complexes. In fact, we show in Corollary 2.22 that one can build a projection complex with respect to $\widehat{X}^{(0)}$, that is, with respect to the union of points in X and translates of Y_i. We also require several additional assumptions to hold, namely Hypothesis 2.19 and Hypothesis 3.1, which ensure that the subgroups $H_1, \dots, H_k$ form a *"sufficiently" spinning family* with respect to the action of G on $\widehat{X}$. Spinning families are defined precisely in Section 3, but can be thought of as a geometric/dynamical generalization of satisfying a small cancellation condition.

Clay and Mangahas [CM22] and Clay, Mangahas, and Margalit [CMM21] studied spinning families in actions on projection complexes. In particular, Clay and Mangahas showed (among other things) that if a collection of subgroups $H_i \leqslant G$ form an equivariant spinning family with respect to the action of G on a projection complex $\mathcal{P}$, then the quotient $\mathcal{P}/\langle\langle H_1, \dots, H_k \rangle\rangle$ is hyperbolic [CM22, Theorem 1.1]. In a similar fashion, in Section 3 we prove hyperbolicity of $\widehat{X}/N$ and establish a criterion for acylindrical hyperbolicity of the quotient; see Theorem 3.13 and Corollary 3.16, respectively. The core of the proof of Theorem 3.13 is to produce a closed lift $T \subseteq \widehat{X}$ of a given geodesic triangle $\overline{T} \subseteq \widehat{X}/N$. Since $\widehat{X}$ is hyperbolic, the triangle T is uniformly slim, and therefore so is $\overline{T}$, as the quotient map $\widehat{X} \to \widehat{X}/N$ is 1–Lipschitz. In turn, in order to find a closed lift of $\overline{T}$, we start with an open lift of the triangle $\overline{T}$ and define a "bending procedure" that eventually yields a closed lift; see Lemma 3.11 and the surrounding discussion. The main ingredient in the above procedure is the existence of a *shortening pair* (Proposition 3.8), which can be thought of as a version of Greendlinger's Lemma for a spinning family [Gre60, CM22]. An analogous procedure shows that we can also lift quadrangles from $\widehat{X}/N$ to $\widehat{X}$, which we use to prove that acylindrical hyperbolicity is preserved under taking quotients by sufficiently spinning families. Similar ideas of lifting polygons from quotients by collections of subgroups satisfying similar small-cancellation-like conditions appear in [Dah18, DHS21, BHMS24, MS25, MS24, CM22]. Indeed, our arguments follow similar lines as those of Clay–Mangahas, with necessary adjustments to handle the fact that $\widehat{X}$ itself is not a projection complex.

In Section 4, we focus on the case that G is a hierarchically hyperbolic group. Hierarchically hyperbolic groups are defined precisely in Definition 4.7, but we describe a few key aspects here. Roughly, an HHG structure on a group G consists of a collection of projections from G onto hyperbolic spaces $\{\mathcal{C}U \mid U \in \mathfrak{S}\}$ indexed by a set $\mathfrak{S}$. There are three relations on $\mathfrak{S}$: nesting, transversality, and orthogonality. Intuitively, the projections onto hyperbolic spaces encode "hyperbolic pieces" of G, while the relations encode how the various pieces fit together to build the entire group. The most relevant relation for this outline is nesting, which is a partial order on $\mathfrak{S}$ with a unique largest element, typically denoted by S and called the *top-level domain*. The hyperbolic space associated with the top-level domain is called the *top-level space*.

The main technical result of Section 4 is Theorem 4.13, which states that the quotient of a HHG $(G, \mathfrak{S})$ by the normal closure N of a sufficiently spinning family of subgroups $\{H_1, \ldots, H_k\}$ is a HHG. We present an informal version of this result and a brief description of the HHG structure of the quotient here; see Construction 4.26 for further details on the structure.

Theorem E. *Let $(G, \mathfrak{S})$ be a (relative) HHG with top-level coordinate space X. Suppose there exist quasiconvex, geometrically separated subspaces $Y_1, \ldots, Y_k \subseteq X$ with metrically proper, cobounded actions of subgroups $H_1, \ldots, H_k$. If the subgroups $H_1, \ldots, H_k$ form a sufficiently spinning family with respect to a certain cone-off of X, then the quotient of G by $N = \langle\langle H_1, \ldots, H_k \rangle\rangle$ is a (relative) HHG, with the following structure.*

- *The set of domains is $\mathfrak{S}/N = \{\overline{U} \mid U \in \mathfrak{S}\}$.*
- *If $\overline{U} \neq \overline{S}$, then the associated hyperbolic space is isometric to $\mathcal{C}U$ for some (equivalently, any) representative $U \in \mathfrak{S}$ of $\overline{U}$.*
- *The top-level domain is $\widehat{X}/N$, where $\widehat{X}$ is the cone-off of X with respect to the translates of $\{Y_1, \ldots, Y_k\}$. Furthermore, $\widehat{X}/N$ is quasi-isometric to X/N.*
- *Two domains $\overline{U}, \overline{V} \in \mathfrak{S}/N$ are orthogonal or nested if they admit orthogonal or nested representatives, respectively, and are transverse otherwise.*

The key idea in showing that the above candidate structure satisfies the axioms of an HHG is to define preferred representatives of elements of G/N, $\widehat{X}/N$, and $\mathfrak{S}/N$ in $G, \widehat{X}$, and $\mathfrak{S}$, respectively. For this, we introduce the notion of *minimal representatives* in Definition 3.10. Briefly, if $\overline{x}, \overline{y} \in \widehat{X}/N$, for example, then representatives x, y of $\overline{x}, \overline{y}$, respectively, are minimal if $\mathsf{d}_{\widehat{X}}(x,y) = \mathsf{d}_{\widehat{X}/N}(\overline{x}, \overline{y})$. For domains, the rough idea is that $U, V \in \mathfrak{S}$ are minimal representatives if the distance between the images of ρ^U_S and ρ^V_S in $\widehat{X}$ is minimal among all possible representatives. The relationship (i.e., nesting, transversality, or orthogonality) between domains in $\mathfrak{S}/N$ is then defined to be the relationship between minimal representatives of those domains in $\mathfrak{S}$, and we similarly use minimal representatives to define the various projections. To verify each axiom, the general strategy is to use minimal representatives to 'lift' the setup of the axiom from $(G/N, \mathfrak{S}/N)$ to $(G, \mathfrak{S})$, where the axiom is satisfied, and then push the result back down to $(G/N, \mathfrak{S}/N)$. The technicalities involved in "lifting" the setup of the axioms are dealt with in Section 4.4 and Section 4.5, and the axioms themselves are verified in Section 4.6.

Finally, in Section 5, we study random walks on an acylindrically hyperbolic group. We show that subgroups generated by finitely many independent random elements a.a.s. satisfy Hypothesis 2.19, Hypothesis 3.1, and the assumptions of Corollary 3.16, thus proving Theorem A. Under then additional assumption that the group is an HHG, we show that such subgroups a.a.s. satisfy the assumptions of Theorem 4.13, and this is then used to prove Theorem B, Corollary C, and Corollary D in Section 5.3.

Comparison with very rotating quotients

Theorem 4.13 should be compared to [BHS17a, Theorem 6.1], which implies the same result for a *very rotating family* of subgroups, rather than a sufficiently spinning family. While spinning families and rotating families capture similar behavior, it is not clear that [BHS17a, Theorem 6.1] can be used to prove that random quotients of HHGs are HHGs. The key difference in the construction of a HHG structure on the quotient in each case is the relationship between the various constants involved.

In [BHS17a], the first step is to modify the top-level hyperbolic space $\mathcal{C}S$ by gluing on "hyperbolic cones" as in [DGO17] to obtain a new hyperbolic space $\widehat{\mathcal{C}S}$. The collection of random subgroups $\langle w_{i,n} \rangle$ would then have to be an r–rotating family with respect to $\widehat{\mathcal{C}S}$ for sufficiently large r. The problem with using this construction is controlling the growth of r as n tends to infinity. Precisely how large the constant r needs to be is not completely clear, but a careful reading of the proofs in [BHS17a] shows that it depends at least linearly on the geometric separation constant of the hyperbolic cones. Roughly, to be an r–rotating family, the translation length of the random walks must be sufficiently large with respect to r, and, in fact, exponential in r [DGO17, Theorem 6.35]. Work of Maher–Tiozzo [MT18, MT21] and Maher–Sisto [MS19] show that the geometric separation constant of the random walk grows linearly in n, as discussed in Section 5.1, and so the translation length of the random walk would need to grow *exponentially* in n to be able to use this construction. However, the translation length grows only linearly in n [MT18].

If one could improve the geometric separation constants of the random walk to grow logarithmically in n, it might be possible to use [BHS17a, Theorem 6.1] to obtain the results in this paper, though one would still need a better understanding of the precise relationship between r and the geometric separation constant. Instead, in this paper we use the more straight-forward cone-off procedure described above and replace rotating families with spinning families. With this construction, we show that the collection of random subgroups $\langle w_{i,n} \rangle$ needs to to be an L–spinning family for a constant L that is again linear in the geometric separation constant of the random walks; see Remark 3.2. To form an L–spinning family, however, we only need the translation length to be linear in L, which holds by Theorem 5.12.

Since the precise constants are critical in this paper, we have included a summary of their definitions and relative dependencies in Appendix A. We suggest keeping it handy while reading through the paper.

Acknowledgements.

We thank Alex Sisto for helpful conversations, and particularly for pointing us toward spinning families. We thank Yair Minsky for asking the question about random walks on mapping class groups which was the genesis of this project. The third author is grateful to the other authors for the opportunity to join this project. We also thank Inhyeok Choi for comments on an early draft and encouragement to include Theorem A.

The first author is partially supported by NSF grants DMS-2106906 and DMS-2340341. The second author was supported by the Additional Funding Programme for Mathematical Sciences, delivered by EPSRC (EP/V521917/1) and the Heilbronn Institute for Mathematical Research. The fourth author was partially supported by ISF grant 660/20, and by a Zuckerman Fellowship at the Technion. The fifth author was partially supported by NSF grants DMS-2202986 and DMS-1840190.

2 Background

2.1 Hyperbolic spaces

In this section we review some basic properties of metric spaces, including hyperbolicity. More details can be found in [BH99]. Most lemmas in this section are standard, but we provide proofs for completeness and to make explicit all constants, as they will play an important role later in the paper.

In what follows we consider subsets of a fixed space X with metric $\mathsf{d} = \mathsf{d}_X$. Every metric space (possibly equipped with an action of a group G by isometries) is (G-equivariantly) quasi-isometric to a simplicial graph, by e.g. [CdlH16, Lemma 3.B.6], so we can and will assume that all metric spaces we consider are simplicial graphs.

Given a subspace $A \subseteq X$, the closed R-neighborhood of A in X is denoted by $\mathcal{N}_R^X(A)$, or simply $\mathcal{N}_R(A)$ when the ambient space is understood. Let $\lambda \geqslant 1$ and $c \geqslant 0$. A map $f\colon (X, \mathsf{d}_X) \to (Y, \mathsf{d}_Y)$ of metric spaces is a (λ, c)*–quasi-isometric embedding* if for all $x, y \in X$, we have that

$$\lambda^{-1}\mathsf{d}_X(x,y) - c \leqslant \mathsf{d}_Y(f(x), f(y)) \leqslant \lambda \mathsf{d}_X(x,y) + c.$$

A *quasi-isometry* is a quasi-isometric embedding which is also *coarsely surjective*, meaning that $Y \subseteq \mathcal{N}_R(f(X))$ for some constant $R \geqslant 0$.

A (λ, c)*–quasigeodesic* in X is a (λ, c)–quasi-isometric embedding of an interval into X. When the constants λ and c are the same, we simply call such path a λ-quasigeodesic. A *geodesic* is a $(1,0)$-quasigeodesic, that is, an isometric embedding of an interval. Given points $x, y \in X$, we denote a geodesic from x to y by $[x, y]^X$; if the space X is clear from context we simply write $[x, y]$.

A metric space is *geodesic* (resp. (λ, c)-quasigeodesic) if any two points are connected by a geodesic (resp. (λ, c)-quasigeodesic). For $\delta \geqslant 0$, a geodesic metric space X is δ*–hyperbolic* if, for every three points $x, y, z \in X$, we have $[x, y] \subseteq \mathcal{N}_\delta([x, z] \cup [z, y])$; we say that geodesic triangles in a δ-hyperbolic space are δ*-slim*. If the particular choice of δ is not important, we simply say that X is *hyperbolic*. The *boundary* ∂X of a hyperbolic space is the set of quasigeodesic rays $[0, \infty) \to X$ up to bounded Hausdorff distance (see e.g. [BH99, III.H.3]).

In this paper, all quasigeodesics $\gamma\colon I \to X$ will be continuous. In a hyperbolic space, this is no loss of generality by [BH99, Lemma III.H.1.11]. We denote the *length* of the quasigeodesic γ in X by $\ell_X(\gamma)$.

Hyperbolic spaces satisfy the following *Morse* property, which states that quasigeodesics with the same endpoints remain in a uniform neighborhood of each other (see, e.g., [BH99, III.H.1.7]). This is also known as *quasigeodesic stability*.

Lemma 2.1. *For all $\delta, c \geqslant 0$, $\lambda \geqslant 1$, there is a constant $\Phi = \Phi(\lambda, c, \delta) \geqslant 0$ satisfying the following. Let X be a δ–hyperbolic space, and let γ_1, γ_2 be (λ, c)–quasigeodesics with the same endpoints in $X \cup \partial X$. Then the Hausdorff distance between γ_1 and γ_2 is at most Φ.*

A subspace Y of a hyperbolic space X is K*-quasiconvex* if for any $x, y \in Y$ we have $[x, y] \subset \mathcal{N}_K(Y)$.

Lemma 2.2. *Let X be δ-hyperbolic, and let $Z \subseteq X$ be a subset of diameter E. For every $A \geqslant 0$, the neighborhood $\mathcal{N}_A(Z)$ is $(2\delta + E)$-quasiconvex.*

Proof. Let $x, y \in \mathcal{N}_A(Z)$, and let $x', y' \in Z$ be such that $\mathsf{d}(x, x') \leqslant A$ and $\mathsf{d}(y, y') \leqslant A$. From the definition of hyperbolicity, it readily follows that geodesic quadrangles in X are 2δ-slim.

Hence any geodesic $[x, y]$ is in the 2δ-neighborhood of geodesics $[x, x'] \cup [x', y'] \cup [y', y]$. Since $\mathsf{diam}(Z) \leqslant E$, $[x, y]$ is actually in the $(2\delta + E)$-neighborhood of $[x, x'] \cup [y', y]$, and these two geodesics lie in $\mathcal{N}_A(Z)$. ☐

From now on, we assume that X is a δ-hyperbolic *graph*, and every point in X is thought of as a vertex, so that distances between points are integer-valued. If $Y \subseteq X$ is quasiconvex, then for every $z \in X$ one can define a coarse closest point projection $\pi_Y : X \to Y$ by mapping every $z \in X$ to the collection

$$\pi_Y(z) = \{y \in Y \mid \mathsf{d}(z, Y) = \mathsf{d}(z, y)\}. \tag{1}$$

For any such $Y \subseteq X$ and every $A, B \subseteq X$ we set $\mathsf{d}^{\pi}_Y(A, B) = \mathsf{diam}(\pi_Y(A) \cup \pi_Y(B))$.

Lemma 2.3 (Closest point projections are uniformly Lipschitz)**.** *Let X be a δ–hyperbolic graph, $Y \subseteq X$ a K–quasiconvex subspace, and $\pi_Y : X \to Y$ a coarse closest point projection. Let $x, y \in X$ be such that $\mathsf{d}(x, y) \leqslant 1$. Then $\mathsf{d}^{\pi}_Y(x, y) \leqslant J \coloneqq 2K + 10\delta + 2$. In particular, for any $w, z \in X$,*

$$\mathsf{d}^{\pi}_Y(w, z) \leqslant J \max\{\mathsf{d}_X(w, z), 1\}.$$

Proof. Let $x' \in \pi_Y(x)$ and $y' \in \pi_Y(y)$, and consider a geodesic quadrangle with vertices $\{x, x', y', y\}$. Towards a contradiction, if $\mathsf{d}(x', y') > J$, we can find a point $z \in [x', y']$ such that both $\mathsf{d}(x', z)$ and $\mathsf{d}(y', z)$ are at least $K + 5\delta$. Let $z' \in Y$ be such that $\mathsf{d}(z, z') \leqslant K$, which exists as Y is K–quasiconvex. Since quadrangles in X are 2δ–slim, there exists $w \in [x', x] \cup [x, y] \cup [y, y']$ within distance 2δ from z. If $w \in [x', x]$ then

$$\mathsf{d}(x, x') = \mathsf{d}(x, w) + \mathsf{d}(w, x') \geqslant \mathsf{d}(x, w) + \mathsf{d}(z, x') - \mathsf{d}(w, z) \geqslant \mathsf{d}(x, w) + K + 3\delta,$$

while

$$\mathsf{d}(x, z') \leqslant \mathsf{d}(x, w) + \mathsf{d}(w, z') \leqslant \mathsf{d}(x, w) + K + 2\delta.$$

This would contradict the definition of x' as a closest point to x in Y. For the same reason, w cannot lie on $[y, y']$, and so $w \in [x, y] - \{x, y\}$. However, this is impossible, as $\mathsf{d}(x, y) \leqslant 1$. ☐

Lemma 2.4 ([CDP90, Proposition 10.2.1])**.** *Given $\delta, K \geqslant 0$, there is a constant $\Omega = \Omega(\delta, K)$ such that the following holds. Let Y be a K-quasiconvex subset of a δ-hyperbolic graph X. For any pair of points $x, y \in X$, any $x' \in \pi_Y(x)$, and any $y' \in \pi_Y(y)$, if $\mathsf{d}(x', y') \geqslant \Omega$, then the* nearest point path *$[x, x'] \cup [x', y'] \cup [y', y]$ is a $(1, \Omega)$-quasigeodesic.*

A collection $\mathcal{Y}$ of quasiconvex subspaces of X is *geometrically separated* if for every $\varepsilon > 0$, there exists $R > 0$ such that $\mathsf{diam}(\mathcal{N}_\varepsilon(Y') \cap Y) \leqslant R$ whenever $Y \neq Y' \in \mathcal{Y}$.

Lemma 2.5. *Let $\mathcal{Y}$ be a collection of K-quasiconvex subspaces of the δ-hyperbolic graph X. Suppose there exists $M_0 > 0$ such that $\mathsf{diam}(\mathcal{N}_{2K+2\delta}(Y') \cap Y) \leqslant M_0$ whenever $Y, Y' \in \mathcal{Y}$ are distinct. Then $\mathcal{Y}$ is geometrically separated. More precisely, for every $t \geqslant 0$,*

$$\mathsf{diam}(\mathcal{N}_t(Y') \cap Y) \leqslant M(t) \coloneqq M_0 + 2K + 2t + 4\delta + 2. \tag{2}$$

If the assumption of Lemma 2.5 holds, we say that the collection $\mathcal{Y}$ is M_0-*geometrically separated.*

Proof of Lemma 2.5. Let $x, y \in \mathcal{N}_t(Y') \cap Y$, and let $x', y' \in Y'$ realize the distance between Y' and x, y, respectively. If $\mathsf{d}(x, y) \leqslant 2t + 4\delta + 1$ there is nothing to prove. Otherwise, consider a geodesic quadrangle $[x, y] \cup [y, y'] \cup [y', x'] \cup [x', x]$, and let $a \in [x, y]$ be such that $t + 2\delta \leqslant \mathsf{d}(x, a) \leqslant t + 2\delta + 1$. Notice that $\mathsf{d}(a, y) \geqslant \mathsf{d}(x, y) - \mathsf{d}(x, a) \geqslant t + 2\delta$ as well. Since this geodesic quadrangle is 2δ-slim, there exists $a' \in [y, y'] \cup [y', x'] \cup [x', x]$ such that $\mathsf{d}(a, a') \leqslant 2\delta$. Notice that a' cannot lie in the interior of $[x, x']$, as otherwise the reverse triangle inequality would give that $\mathsf{d}(x, x') > \mathsf{d}(a', x) \geqslant \mathsf{d}(x, a) - \mathsf{d}(a, a') \geqslant t$. For the same reason, a' does not lie in the interior of $[y, y']$, so we must have that $a' \in [x, y] \subseteq \mathcal{N}_K(Y')$. Hence $a \in \mathcal{N}_{K+2\delta}(Y')$. In turn, as a lies on a geodesic $[x, y]$ with endpoints in Y, there is some $p \in Y$ such that $\mathsf{d}(a, p) \leqslant K$. Thus we have found a point $p \in Y \cap \mathcal{N}_{2K+2\delta}(Y')$ such that $\mathsf{d}(x, p) \leqslant \mathsf{d}(x, a) + \mathsf{d}(a, p) \leqslant t + 2\delta + K + 1$. An analogous argument produces a point $q \in Y \cap \mathcal{N}_{2K+2\delta}(Y')$ such that $\mathsf{d}(y, q) \leqslant t + 2\delta + K + 1$. Since $p, q \in \mathcal{N}_{2K+2\delta}(Y') \cap Y$, which has diameter at most M_0, we have

$$\mathsf{d}(x, y) \leqslant \mathsf{d}(x, p) + \mathsf{d}(p, q) + \mathsf{d}(q, y) \leqslant M_0 + 2K + 2t + 4\delta + 2. \qquad \square$$

From now on, we shall say that a function f, depending on some constants $c_1, \ldots, c_n$, and M_0, is *bounded linearly* in M_0 if there are positive functions $a(c_1, \ldots, c_n)$ and $b(c_1, \ldots, c_n)$ such that $|f(M_0, c_1, \ldots, c_n)| \leqslant a(c_1, \ldots, c_n)M_0 + b(c_1, \ldots, c_n)$.

Lemma 2.6. *Let X be a δ-hyperbolic graph, and let $\mathcal{Y}$ be an M_0-geometrically separated collection of K-quasiconvex subspaces. There exists $B = B(\delta, K, M_0)$ which is bounded linearly in M_0 and such that $\mathsf{diam}\pi_{Y'}(Y) \leqslant B$ for every $Y \neq Y' \in \mathcal{Y}$.*

Proof. We claim that it suffices to take $B = M(2K + 7\delta + 1)$, where M is the geometric separation function from (2), which is bounded linearly in M_0. Let $x, y \in Y$, and let $x' \in \pi_{Y'}(x)$ and $y' \in \pi_{Y'}(y)$. If $\mathsf{d}(x', y') \leqslant 2K + 10\delta + 2$ we are done, as the latter is less than $M(2K + 7\delta + 1)$. Otherwise consider a geodesic quadrangle with vertices $\{x, x', y', y\}$. Let $z' \in [x', y']$ be such that $K + 5\delta \leqslant \mathsf{d}(x', z') \leqslant K + 5\delta + 1$. By slimness of quadrangles, one can find a point $w \in [x', x] \cup [x, y] \cup [y, y']$ within distance 2δ from z. Arguing exactly as in Lemma 2.3, one sees that w must belong to $[x, y]$, or it would violate the fact that x' (resp. y') realizes the distance between Y' and x (resp. y). Therefore

$$\mathsf{d}(x', Y) \leqslant \mathsf{d}(x', z) + \mathsf{d}(z, w) + \mathsf{d}(w, Y) \leqslant 2K + 7\delta + 1,$$

where we used that Y is K-quasiconvex to bound $\mathsf{d}(w, Y)$. The same argument works for y', so the two points lie in $Y' \cap \mathcal{N}_{2K+7\delta+1}(Y)$. By Lemma 2.5, the diameter of the latter is bounded by $M(2K + 7\delta + 1)$, as required. $\square$

2.2 Cone-offs of graphs

Let $\mathcal{Y}$ be a collection of subgraphs of a connected hyperbolic graph X. We denote by $\widehat{X}$ the graph obtained from X as follows: add a vertex v_Y for each $Y \in \mathcal{Y}$, and add edges from v_Y to each vertex in Y. Note that X is naturally a subspace of $\widehat{X}$. We say that $\widehat{X}$ is the *cone-off*, or *electrification*, of X with respect to $\mathcal{Y}$, and we call the vertices v_Y *cone vertices*.

In what follows, the endpoints of a path $\gamma \colon [0, 1] \to X$ are denoted by $\gamma_- = \gamma(0)$ and $\gamma_+ = \gamma(1)$, and $\alpha * \beta$ denotes the concatenation of two paths α and β such that $\alpha_+ = \beta_-$.

Definition 2.7. Let $\widehat{X}$ be the cone-off of X with respect to a family of subgraphs $\mathcal{Y}$. Let $\gamma = u_1 * e_1 * \cdots * e_n * u_{n+1}$ be a concatenation of geodesic segments where each e_i is a concatenation of two edges sharing a common cone vertex and the u_i are (possibly trivial) segments contained in X. A *de-electrification* $\widetilde{\gamma}$ of γ is the concatenation $u_1 * \eta_1 * \cdots * \eta_n * u_{n+1}$, where each η_i is a geodesic in X connecting the endpoints of e_i. If e_i connects points of $Y \in \mathcal{Y}$, then η_i is a *Y-component of* $\widetilde{\gamma}$.

We are particularly interested in the case that $\mathcal{Y}$ is a collection of uniformly quasiconvex subgraphs.

Lemma 2.8 ([KR14, Proposition 2.6])**.** *Let X be a δ-hyperbolic graph, and let $\mathcal{Y}$ be a collection of K-quasiconvex subspaces. There exists $\widehat{\delta} = \widehat{\delta}(\delta, K)$ such that $\widehat{X}$ is $\widehat{\delta}$-hyperbolic.*

Moreover, de-electrifications of geodesics in $\widehat{X}$ are uniformly close to geodesics in X.

Lemma 2.9 ([Spr18, Corollary 2.23])**.** *Let X be a δ-hyperbolic graph, let $\mathcal{Y}$ be a collection of K-quasiconvex subsets of X, and let $\widehat{X}$ be the cone-off of X with respect to $\mathcal{Y}$. Then there exists a constant $D = D(\delta, K)$ such that, for any pair of points $x, y \in X$, every geodesic $[x, y]^X$, and every geodesic γ from x to y in $\widehat{X}$, we have*

$$[x, y] \subseteq \mathcal{N}_D^X(\widetilde{\gamma}),$$

where $\widetilde{\gamma}$ is a de-electrification of γ.

Remark 2.10. [Spr18, Corollary 2.23] is stated for *connected* subgraphs; however, the latter hypothesis is only used to ensure that de-electrifications of geodesics exist, which is true under the requirement that X is hyperbolic (in particular geodesic, hence path-connected). Moreover, the cone-off X' of X used in [Spr18] is slightly different: there, an edge is added between any two vertices lying in a common $Y \in \mathcal{Y}$. There is a map from X' to $\widehat{X}$ sending an edge to a concatenation of two edges between the same pair of points, and it is immediate that de-electrifications $\widetilde{\gamma}$ of a geodesic in X' agree with de-electrifications of the image of γ in $\widehat{X}$. Hence Lemma 2.9 still holds for $\widehat{X}$.

Lemma 2.9 leads to the following corollary.

Corollary 2.11. *In the setting of Lemma 2.9, let $x, y, w \in X$, and let γ be an $\widehat{X}$–geodesic from x to y. For any $t \geqslant 0$, if $\mathsf{d}_X(w, [x, y]^X) \leqslant t$, then $\mathsf{d}_{\widehat{X}}(w, \gamma) \leqslant t + D + K + 1$.*

Proof. Let $\widetilde{\gamma}$ be the de-electrification of γ. Let $w' \in [x, y]^X$ satisfy $\mathsf{d}_X(w, w') \leqslant t$. By Lemma 2.9, there is a point $z \in \widetilde{\gamma}$ with $\mathsf{d}_X(w', z) \leqslant D$. Either z lies on γ, in which case we have proven the bound, or z lies on a Y–component of $\widetilde{\gamma}$ for some $Y \in \mathcal{Y}$. In the latter case, the K–quasiconvexity of Y implies that there exists $z' \in Y$ such that $\mathsf{d}_X(z, z') \leqslant K$, and there is an edge from z' to v_Y in $\widehat{X}$. The cone vertex v_Y lies on γ by construction, and so $\mathsf{d}_{\widehat{X}}(w, \gamma) \leqslant t + D + K + 1$, as desired. □

Definition 2.12. For every $Y \in \mathcal{Y}$ there is a set-valued projection $\widehat{X} - \{v_Y\} \to 2^Y$, which we still denote by π_Y, defined as follows. For every $x \in X$ the projection is $\pi_Y(x)$, and for every $U \in \mathcal{Y}$ other than Y we set $\pi_Y(v_U) \coloneqq \pi_Y(U)$, where $\pi_Y(U)$ is as defined in (1), considering U as a subspace of X. For every $x, y \in \widehat{X} - \{v_Y\}$, we set $\mathsf{d}_Y^\pi(x, y) = \mathsf{diam}(\pi_Y(x) \cup \pi_Y(y))$.

The following strengthening of the bounded geodesic image property describes how cone points relate to geodesics joining a pair of points.

Lemma 2.13 (Strong bounded geodesic image)**.** *In the setting of Lemma 2.9, suppose further that the family $\mathcal{Y}$ is M_0-geometrically separated, in the sense of Lemma 2.5. There exists a constant $C = C(\delta, K, M_0)$ bounded linearly in M_0 such that for every $Y \in \mathcal{Y}$ and $x, y \in \widehat{X} - \{v_Y\}$, if a $\widehat{X}$–geodesic γ does not pass through v_Y then $\mathsf{d}^\pi_Y(x, y) \leqslant C$.*

Proof. Assume first that $x, y \in X$. Set $\mathcal{R} = 2\delta + 2K + D + 2$, where $D = D(\delta, K)$ is the constant from Lemma 2.9. Our first goal is to prove that if $\mathsf{d}^\pi_Y(x, y) > J$, where $J = J(\delta, K)$ is the constant from Lemma 2.3, then γ intersects the $\mathcal{R}$-neighborhood of v_Y in $\widehat{X}$. To that end, let $[x, y]$ be an X-geodesic between x and y. Since $\mathsf{d}^\pi_Y(x, y) > J$, we can argue as in Lemma 2.3 to find some $z' \in Y$ such that $\mathsf{d}_X(z', [x, y]) \leqslant K + 2\delta$. Corollary 2.11 then gives

$$\mathsf{d}_{\widehat{X}}(v_Y, \gamma) \leqslant 1 + \mathsf{d}_{\widehat{X}}(z', \gamma) \leqslant 1 + \mathsf{d}_X(z', [x, y]) + D + K + 1 \leqslant 2\delta + 2K + D + 2 = \mathcal{R}.$$

Next, let a, b be the first and last vertices of $\gamma \cap X$ contained in the $(\mathcal{R} + 2)$-neighborhood of v_Y. Notice that $\mathsf{d}_{\widehat{X}}(a, v_Y) \geqslant \mathcal{R} + 1$, and similarly for b, since no two points on $\gamma - X$ are adjacent. Hence by the above argument both $\mathsf{d}^\pi_Y(x, a)$ and $\mathsf{d}^\pi_Y(b, y)$ are at most J. Moreover, if we let $a = a_0, a_1, \ldots, a_n = b$ be the subsegment of γ between a and b, then $n \leqslant \mathsf{d}_{\widehat{X}}(a, b) \leqslant 2\mathcal{R} + 4$. Finally, for every $0 \leqslant i \leqslant n - 1$, we have $\mathsf{d}^\pi_Y(a_i, a_{i+1}) \leqslant \max\{J, B\}$, where $B = B(\delta, K, M)$ is the constant from Lemma 2.6. Indeed, either a_i, a_{i+1} are X-adjacent, and therefore $\mathsf{d}^\pi_Y(a_i, a_{i+1}) \leqslant J$, or $a_i = v_U$ for some $U \neq Y$ and $a_{i+1} \in U$ (or vice versa), so $\mathsf{diam}(\pi_Y(a_{i+1}) \cup \pi_Y(a_i)) \leqslant \mathsf{diam}(\pi_Y(U)) \leqslant B$. The triangle inequality thus yields

$$\mathsf{d}^\pi_Y(x, y) \leqslant \mathsf{d}^\pi_Y(x, a) + \sum_{i=0}^{n-1} \mathsf{d}^\pi_Y(a_i, a_{i+1}) + \mathsf{d}^\pi_Y(b, y) < C_0 \coloneqq 2J + (2\mathcal{R} + 4)\max\{B, J\},$$

concluding the proof of Lemma 2.13 for points in X.

Finally, let $x, y \in \widehat{X}$ be any two points, let γ be an $\widehat{X}$-geodesic between them which does not pass through v_Y, and let x' and y' be the first and last points of $\gamma \cap X$, respectively. In particular, either $x = x'$ or $x = v_U$ for some $U \in \mathcal{Y}$ containing x', and similarly for y. Note that $\mathsf{d}^\pi_Y(x', x)$, $\mathsf{d}^\pi_Y(y, y') \leqslant B$, and $\mathsf{d}^\pi_Y(x', y') \leqslant C_0$ by the above argument, so the triangle inequality yields that $\mathsf{d}^\pi_Y(x, y) \leqslant C \coloneqq C_0 + 2B$. The constant C is bounded linearly in M_0, as so are C_0 and B. ☐

2.3 Projection complexes

In this section, we recall the machinery of projection complexes, first introduced by Bestvina, Bromberg, and Fujiwara in [BBF15]. Let $\mathcal{Y}$ be a collection of metric spaces. Suppose that for each $Y \in \mathcal{Y}$ there is a projection π_Y from the elements of $\mathcal{Y} - \{Y\}$ to subsets of Y. Then we may define a "distance function" $\mathsf{d}^\pi_Y : (\mathcal{Y} - \{Y\})^2 \to [0, \infty)$ by

$$\mathsf{d}^\pi_Y(U, V) = \mathsf{diam}\left(\pi_Y(U) \cup \pi_Y(V)\right).$$

(Note that this is precisely what we did in Definition 2.12). The map d^π_Y is usually not a true distance function since we may have $\mathsf{d}^\pi_Y(U, U) > 0$.

Definition 2.14 ((Strong) projection axioms)**.** Given a collection $\mathcal{Y}$ and distance functions $\{\mathsf{d}^\pi_Y\}_{Y \in \mathcal{Y}}$ as above, as well as a constant $\theta \geqslant 0$, the *projection axioms for $\mathcal{Y}$ with constant θ* are

(I) $\mathsf{d}^\pi_Y(U, V) = \mathsf{d}^\pi_Y(V, U)$;

(II) $\mathsf{d}^{\pi}_{Y}(U, W) \leqslant \mathsf{d}^{\pi}_{Y}(U, V) + \mathsf{d}^{\pi}_{Y}(V, W)$.

(III) $\mathsf{d}^{\pi}_{Y}(U, U) \leqslant \theta$;

(IV) $\min\{\mathsf{d}^{\pi}_{Y}(U, V), \mathsf{d}^{\pi}_{V}(U, Y)\} \leqslant \theta$; and

(V) for all $U, V \in \mathcal{Y}$, the set $\{Y \in \mathcal{Y} \mid \mathsf{d}^{\pi}_{Y}(U, V) \geqslant \theta\}$ is finite.

The *strong projection axioms for $\mathcal{Y}$ with constant θ* were defined in [BBFS19], and are obtained by replacing Item (IV) by the following stronger statement:

(IV') If $\mathsf{d}^{\pi}_{Y}(U, V) > \theta$ then $\mathsf{d}^{\pi}_{V}(U, W) = \mathsf{d}^{\pi}_{V}(Y, W)$ for all $W \in \mathcal{Y} - \{V\}$.

Bestvina–Bromberg–Fujiwara–Sisto show that functions d^{π}_{Y} satisfying the projection axioms can be modified to satisfy the strong projection axioms:

Theorem 2.15 ([BBFS19, Theorem 4.1])**.** *Assume that $\mathcal{Y}$ is a collection of metric spaces together with functions $\{\mathsf{d}^{\pi}_{Y}\}_{Y \in \mathcal{Y}}$ satisfying the projection axioms with constant θ. Then there are functions $\mathsf{d}_{Y} \colon (\mathcal{Y} - \{Y\})^2 \to [0, \infty)$ satisfying the strong projection axioms with constant 11θ, and such that for all $Y \in \mathcal{Y}$,*

$$\mathsf{d}^{\pi}_{Y} - 2\theta \leqslant \mathsf{d}_{Y} \leqslant \mathsf{d}^{\pi}_{Y} + 2\theta.$$

Definition 2.16. Suppose that $(\mathcal{Y}, \{\mathsf{d}^{\pi}_{Y}\}_{Y \in \mathcal{Y}})$ satisfy the projection axioms with constant θ, and let $\{\mathsf{d}_{Y}\}_{Y \in \mathcal{Y}}$ be the modified distance functions. For any $\mathfrak{K} \geqslant 0$, the *projection complex* $\mathcal{P}_{\mathfrak{K}}(\mathcal{Y})$ is defined as follows. The vertices of $\mathcal{P}_{\mathfrak{K}}(\mathcal{Y})$ are $\mathcal{Y}$, and two vertices $U, V \in \mathcal{Y}$ are joined by an edge if $\mathsf{d}_{Y}(U, V) \leqslant \mathfrak{K}$ for all $Y \in \mathcal{Y} - \{U, V\}$. When the collection $\mathcal{Y}$ is unimportant or clear from context we use the notation $\mathcal{P}_{\mathfrak{K}}$.

We recall some facts about projection complexes.

Lemma 2.17 (Bounded path image, [BBFS19, Corollary 3.4])**.** *If $\mathfrak{K} \geqslant 33\theta$ and a path $U_1, \ldots, U_k$ in $\mathcal{P}_{\mathfrak{K}}$ does not intersect the 2-neighborhood of a vertex Y, then $\mathsf{d}_{Y}(U_1, U_k) \leqslant 11\theta$.*

Corollary 2.18 (Strong bounded geodesic image)**.** *If $\mathfrak{K} \geqslant 33\theta$ and a geodesic $U_1, \ldots, U_k$ in $\mathcal{P}_{\mathfrak{K}}$ does not contain a vertex Y, then $\mathsf{d}_{Y}(U_1, U_k) \leqslant 22\theta + 6\mathfrak{K}$.*

Proof. Let U_i, U_j be the first and last point of the geodesic within distance 3 from Y. Then $j - i = \mathsf{d}(U_i, U_j) \leqslant 6$. Thus

$$\mathsf{d}_{Y}(U_1, U_k) \leqslant \mathsf{d}_{Y}(U_1, U_i) + \mathsf{d}_{Y}(U_i, U_{i+1}) + \ldots + \mathsf{d}_{Y}(U_{j-1}, U_j) + \mathsf{d}_{Y}(U_j, U_k).$$

Both $\mathsf{d}_{Y}(U_1, U_i)$ and $\mathsf{d}_{Y}(U_j, U_k)$ are at most 11θ by Lemma 2.17. Furthermore, for every $i \leqslant n \leqslant j - 1$ we have that $\mathsf{d}_{Y}(U_n, U_{n+1}) \leqslant \mathfrak{K}$ by definition of the distance in $\mathcal{P}_{\mathfrak{K}}(\mathcal{Y})$, and the bound follows. □

2.3.1 A projection complex from a separated family of quasiconvex subspaces

Quasiconvex subsets of hyperbolic spaces naturally give rise to projection complexes (see for example [BBF15]). For the rest of the section, we will work under the following hypothesis.

Hypothesis 2.19. Let X be a connected graph and $\mathcal{Y}$ a collection of subsets of X. Assume there are constants $\delta, K, M_0 > 0$ such that the following hold.

1. X is δ-hyperbolic.
2. Each $Y \in \mathcal{Y}$ is K–quasiconvex in X.
3. $\mathcal{Y}$ is M_0–geometrically separated, in the sense of Lemma 2.5.

When the above are satisfied, we say that $(X, \mathcal{Y})$ satisfies Hypothesis 2.19 with respect to (δ, K, M_0). We omit the constants when they are unimportant or clear from context.

The goal of this subsection is to prove the following:

Proposition 2.20. *Suppose $(X, \mathcal{Y})$ satisfies Hypothesis 2.19 with respect to (δ, K, M_0). There exists $\theta = \theta(\delta, K, M_0)$ such that $(\mathcal{Y}, \{\mathsf{d}^{\pi}_Y\}_{Y \in \mathcal{Y}})$ satisfies the projection complex axioms (Definition 2.14) with constant θ. Moreover, θ is bounded linearly in M_0.*

Proof. We determine lower bounds on θ one item at a time. Item (I) and Item (II) follow immediately from the fact that d^{π}_Y is defined in terms of diameters of projections, and Item (III) holds with $\theta \geqslant B$ by Lemma 2.6.

We now move to Item (IV). Let $\theta_1 = 2B + J(2\delta + K + 1)$, where J is the Lipschitz constant from Lemma 2.3. Let $U, V, W \in \mathcal{Y}$ be such that $\mathsf{d}^{\pi}_W(U, V) > \theta_1$. Fix $a \in U$, and let $b \in \pi_V(a)$, $c_a \in \pi_W(a)$, and $c_b \in \pi_W(b)$. Then $\mathsf{d}(c_a, c_b) \geqslant \mathsf{d}^{\pi}_W(U, V) - \mathsf{diam}\pi_W(U) - \mathsf{diam}\pi_W(V) \geqslant J$. As in Lemma 2.3, there thus exists $w \in [a, b]$ within distance at most $2\delta + K$ from W. In turn $\mathsf{d}(\pi_V(w), \pi_V(W)) \leqslant J(2\delta + K)$ as projections are J-Lipschitz. Notice that $b \in \pi_V(w)$, as w lies on $[a, b]$ and $b \in \pi_V(a)$, so

$$\mathsf{d}(\pi_V(U), \pi_V(W)) \leqslant \mathsf{d}(b, \pi_V(W)) \leqslant \mathsf{diam}(\pi_V(w)) + \mathsf{d}(\pi_V(w), \pi_V(W)) \leqslant J(2\delta + K + 1).$$

Hence

$$\mathsf{d}^{\pi}_V(U, W) \leqslant \mathsf{diam}\pi_V(U) + \mathsf{d}(\pi_V(U), \pi_V(W)) + \mathsf{diam}\pi_V(W) \leqslant 2B + J(2\delta + K + 1) =: \theta_1,$$

so Item (IV) holds if $\theta \geqslant \theta_1$.

Now let $D_0 = 2K + 4\delta + M(2K + 4\delta) + 1$, which again is bounded linearly in M_0, and let $\theta = 3JD_0 + 2B + 2J(3\delta + 1)$, which is greater than θ_1. We are left to prove Item (V), i.e., that for every $U \neq V \in \mathcal{Y}$ the set $\{Y \in \mathcal{Y} \mid \mathsf{d}^{\pi}_Y(U, V) \geqslant \theta\}$ is finite.

Let γ be a geodesic from u to v, where $u \in U$ and $v \in V$. For every Y as above, $\mathsf{d}^{\pi}_Y(u, v) \geqslant \theta - 2B \geqslant 2J(3\delta + 1)$, so let $a, b \in \gamma$ be the last point such that $\mathsf{d}^{\pi}_Y(u, a) \leqslant J(3\delta + 1)$ and the first point such that $\mathsf{d}^{\pi}_Y(b, v) \leqslant J(3\delta + 1)$, respectively. See Figure 1. Let $[a, b]$ be the subsegment of γ between a and b, and let $u' \in \pi_Y(u)$ and $v' \in \pi_Y(v)$. By slimness of quadrangles, every $w \in [a, b]$ is 2δ-close to some point $z \in [u, u'] \cup [u', v'] \cup [v', v]$. If $z \in [u, u']$ then $u' \in \pi_Y(z)$, so $\mathsf{d}^{\pi}_Y(u, w) \leqslant J + \mathsf{d}^{\pi}_Y(z, w) \leqslant J(2\delta + 1)$, contradicting the fact that w is between a and b. For the same reason, z cannot lie on $[v, v']$, so $[a, b] \subseteq \mathcal{N}_{2\delta}([u', v']) \subseteq \mathcal{N}_{K+2\delta}(Y)$. Now,

$$\mathsf{d}^{\pi}_Y(a, b) \geqslant (\mathsf{d}^{\pi}_Y(u, v) - 2J(3\delta + 1)) \geqslant (\theta - 2B - 2J(3\delta + 1)) = 3JD_0 > J,$$

so by Lemma 2.3 we have that $\mathsf{d}(a, b) \geqslant \frac{1}{J}\mathsf{d}^{\pi}_Y(a, b) \geqslant 3D_0$. Thus, if we cover γ by finitely many segments $\gamma_1, \ldots, \gamma_\ell$, each of length D_0, we must have that $\gamma_i \subseteq [a, b] \subseteq \mathcal{N}_{K+2\delta}(Y)$ for some $i \leqslant l$.

It is enough to show that if $\gamma_i \subseteq \mathcal{N}_{K+2\delta}(Y) \cap \mathcal{N}_{K+2\delta}(Y')$ for some $Y, Y' \in \mathcal{Y}$, then $Y = Y'$, because then the set $\{Y \in \mathcal{Y} \mid \mathsf{d}^{\pi}_Y(U, V) \geqslant \theta\}$ will have cardinality at most l. To see this, notice that γ_i has length D_0, so there exist $p, q \in Y \cap N_{K+2\delta}(\gamma_i) \subseteq Y \cap N_{2K+4\delta}(Y')$ such that $\mathsf{d}(p, q) \geqslant D_0 - 2(K + 2\delta) = M(2K + 4\delta) + 1$. Since $\mathcal{Y}$ is M_0-separated, this means that $Y = Y'$, as required. ∎

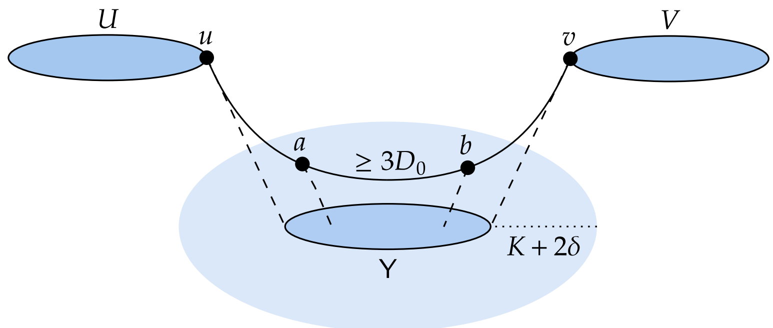

Figure 1: The domains and geodesic involved in the proof of Item (V) in Proposition 2.20. The geodesic $[u, v]$ has a long subsegment in a uniform neighborhood of Y, and the endpoints a and b of this subsegment project to Y close to the endpoints of the whole geodesic.

2.3.2 Adding points to the projection complex

We now define a projection complex structure on $\widehat{X}^{(0)}$ itself, provided that the collection $\mathcal{Y}$ coarsely covers X. That is, we will build the structure using both the points in X and the spaces in $\mathcal{Y}$. This will serve as a unified framework for studying projections between spaces in $\mathcal{Y}$ and from X to spaces in $\mathcal{Y}$.

For every $x \in X$ let $\pi_x\colon X \to \{x\}$ be the constant map, and for every cone point v_Y, let π_{v_Y} be the map $\pi_Y\colon X \to 2^Y$. For every $Y \in \mathcal{Y}$ and $x \in \widehat{X}$, set $\pi_x(v_Y) = \pi_x(Y)$, as in Definition 2.12. Equip the set $\widehat{X}^{(0)}$ with distance functions $\mathsf{d}^\pi_x(y, z) = \mathsf{diam}(\pi_x(y) \cup \pi_x(z))$.

Definition 2.21. Fix a constant $R \geqslant 0$. If $x \in X$, let $\mathcal{V}_R(x) = \{Y \in \mathcal{Y} \mid \mathsf{d}_X(x, Y) \leqslant R\}$ be the collection of *nearest subspaces* to x. If $x = v_Y$ for some $Y \in \mathcal{Y}$, we set $\mathcal{V}_R(x) = \{Y\}$.

Corollary 2.22. *Suppose $(X, \mathcal{Y})$ satisfies Hypothesis 2.19 with respect to (δ, K, M_0), and suppose there exists $R \geqslant 0$ such that $X = \bigcup_{Y\in\mathcal{Y}} \mathcal{N}_R(Y)$. There exists $\Theta = \Theta(\delta, K, M_0, R)$ such that $(\widehat{X}^{(0)}, \{\mathsf{d}^\pi_x\}_{x\in\widehat{X}})$ satisfies the projection complex axioms (Definition 2.14) with constant Θ. Moreover, Θ is bounded linearly in M_0.*

Proof. By Proposition 2.20, $(\mathcal{Y}, \{\mathsf{d}^\pi_Y\}_{Y\in\mathcal{Y}})$ satisfies the axioms with some constant $\theta(\delta, K, M_0)$, which is bounded linearly in M_0. We will show the result holds with $\Theta = \theta + 2(B + JR) + J$.

Item (I) and Item (II) are clear from the construction. Furthermore, for every $x, y \in \widehat{X}$, we have $\mathsf{d}^\pi_y(x, x) \leqslant J \leqslant \Theta$ if $x \in X$, and $\mathsf{d}^\pi_y(x, x) \leqslant \theta \leqslant \Theta$ if $x = v_U$ for some $U \in \mathcal{Y}$, so Item (III) holds.

We next prove Item (IV). If $x, y, z \in \widehat{X}$, then $\min\{\mathsf{d}^\pi_x(y, z), \mathsf{d}^\pi_y(x, z)\} = 0$ unless $x = v_U$ and $y = v_V$ for some $U, V \in \mathcal{Y}$. If $z = v_W$ for some $W \in \mathcal{Y}$, then the bound follows from the corresponding axiom for $(\mathcal{Y}, \{\mathsf{d}^\pi_Y\}_{Y\in\mathcal{Y}})$. Otherwise $z \in X$, and let $Z \in \mathcal{V}_R(z)$. If $Z \in \{U, V\}$, say without loss of generality that $Z = U$, then $\mathsf{d}^\pi_Y(U, z) = \mathsf{d}^\pi_Y(Z, z) \leqslant B + JR$, as π_Y is J-Lipschitz. If instead $Z \notin \{U, V\}$, then

$$\min\{\mathsf{d}^\pi_U(V, z), \mathsf{d}^\pi_V(U, z)\} \leqslant \min\{\mathsf{d}^\pi_U(V, Z), \mathsf{d}^\pi_V(U, Z)\} + B + JR \leqslant \theta + B + JR \leqslant \Theta.$$

For Item (V), let $x, y \in \widehat{X}$, let $U \in \mathcal{V}_R(x)$, and let $V \in \mathcal{V}_R(y)$. If $Z \in \mathcal{Y}$ is such that $\Theta \leqslant \mathsf{d}^\pi_Z(x, y) \leqslant \mathsf{d}^\pi_Z(U, V) + 2(B + JR)$, then $\mathsf{d}^\pi_Z(U, V) \geqslant \theta$, and there are finitely many such Z. This suffices to prove Item (V) as, if $z \in X$, then $\mathsf{d}^\pi_z(x, y) = 0$. ∎

3 Hyperbolic quotients from spinning families

Recall that an isometric action of a group G on a space X is R-*cobounded* for some constant $R \geqslant 0$ if $X = \mathcal{N}_R(G \cdot x)$ for every $x \in X$. In this section, we work under the following standing assumption, which introduces a group action into the framework of Section 2.

Hypothesis 3.1. Let $(X, \mathcal{Y})$ satisfy Hypothesis 2.19 with respect to (δ, K, M_0). Let M be the geometric separation function from Equation (2). Let $J = J(\delta, K)$ be the Lipschitz constant of $\{\pi_Y\}_{Y \in \mathcal{Y}}$ from Lemma 2.3, let $B = B(\delta, K, M_0)$ be the bound on the diameter of projections from Lemma 2.6, and let $C = C(\delta, K, M_0)$ be the bounded geodesic image constant from Lemma 2.13.
Let G be a group acting by isometries on X. Suppose that the following hold.

4. The G-action on X is R-cobounded, for some $R \geqslant 0$; and

5. the collection $\mathcal{Y}$ is G–invariant, that is, for each $Y \in \mathcal{Y}$ every translate gY is in $\mathcal{Y}$.

In particular $X = \bigcup_{Y \in \mathcal{Y}} \mathcal{N}_R(Y)$. Let $\Theta = \Theta(\delta, K, M_0, R)$ be the constant from Corollary 2.22. Set $\widetilde{\Theta} = \max\{\Theta, 2(JR + B + \Theta)/33\}$ and $\text{Ж} = 33\widetilde{\Theta}$, and let $\mathcal{P} \coloneqq \mathcal{P}_{\text{Ж}}(\widehat{X}^{(0)})$ be the associated projection complex, as in Definition 2.16. Let $L_{hyp} = L_{hyp}(\mathcal{P})$ be the threshold from [CM22, Theorem 1.1], and set

$$\overline{L} = \max\{L_{hyp}, 40C, 10(2(B + JR) + 2J + 1)\}. \tag{3}$$

Finally, suppose that the following hold.

6. For each $Y \in \mathcal{Y}$ there is a non-trivial subgroup $H_Y \leqslant \mathrm{Stab}_G(Y)$ such that $gH_Y g^{-1} = H_{gY}$ for each $g \in G$; and

7. There exists $L > \overline{L}$ such that, for any $Y \in \mathcal{Y}$, any $x \neq v_Y \in \widehat{X}$, and any nontrivial $h \in H_Y$, we have $\mathsf{d}^{\pi}_Y(x, hx) > L$.

If the above are satisfied, we say $(X, \mathcal{Y}, G, \{H_Y\}_{Y \in \mathcal{Y}})$ satisfies Hypothesis 3.1 with respect to (δ, K, M_0, R, L). We omit the constants when they are unimportant or clear from context.

We emphasize that the projection complex $\mathcal{P}$ in Hypothesis 3.1 is formed using *all* of the points in $\widehat{X}^{(0)}$, not just the set $\mathcal{Y}$.

Remark 3.2 (Linear dependence of L_{hyp} on M_0)**.** By inspection of [CM22], one sees that the constant L_{hyp} is bounded linearly in M_0. Indeed, Clay and Mangahas first introduce constants C_e, C_p, C_g such that the following hold.

- Whenever x, y are adjacent in $\mathcal{P}$, then $\mathsf{d}^{\pi}_z(x, y) \leqslant C_e$ for every $x \neq z \neq y$. The definition of $\mathcal{P}$ and Theorem 2.15, which states that d^{π}_z and d_z differ by at most 2Θ, ensures that we can choose $C_e = \text{Ж} + 2\Theta$.

- Whenever x, y are joined by a path in $\mathcal{P}$ which does not pass through the 2-neighborhood of z, then $\mathsf{d}^{\pi}_z(x, y) \leqslant C_p$. By Lemma 2.17 we can choose $C_p = 11\widetilde{\Theta} + 2\Theta$.

- Whenever x, y are joined by a geodesic in $\mathcal{P}$ which does not pass through z, then $\mathsf{d}^{\pi}_z(x, y) \leqslant C_g$. By Corollary 2.18 we can use $C_g = 22\widetilde{\Theta} + 6\text{Ж} + 2\Theta$.

From here, they set:

- $m = 11C_e + 6C_g + 5C_p$ and $L_0 = 4(m + \Theta) + 1$ [CM22, Lemma 2.1];
- $L_{short} = \max\{L_0, 5m, 14\Theta\}$ [CM22, Proposition 3.2];
- $L_{lift}(0) = \max\{L_{short}, 40C_g\}$ [CM22, Proposition 4.3];
- $L_{hyp} = L_{lift}(0)$ [CM22, Theorem 1.1].

In the end, L_{hyp} is a piecewise linear function of $\widetilde{\Theta}$ and Θ, both of which are bounded linearly in M_0. We also stress that, as pointed out in [CM22, Lemma 2.1], $L_{hyp} \geqslant L_0$ is also greater than the constant $L(\mathcal{P})$ from [CMM21, Theorem 1.6], so all results of that paper concerning only the action $G \curvearrowright \mathcal{P}$ still hold in our setting.

3.1 Properties of spinning families

For each $x \in \widehat{X}$, define $H_x = H_Y$ if $x = v_Y$ for some $Y \in \mathcal{Y}$, and define $H_x = \{1\}$ if $x \in X$. The collection of subgroups $\{H_x\}_{x \in \widehat{X}}$ form a *spinning family*, a notion introduced by Clay, Mangahas, and Margalit [CMM21, Section 1.7].

Definition 3.3 (Spinning family)**.** Let $\mathcal{P}$ be a projection complex, and let G be a group that acts on $\mathcal{P}$ by isometries. For each vertex $x \in \mathcal{P}$, let H_x be a subgroup of the stabilizer of x in G. Let $L > 0$. The family of subgroups $\{H_x\}$ is an *equivariant L–spinning family* of subgroups of G if it satisfies the following two conditions:

- *Equivariance:* If $g \in G$ and x is a vertex of $\mathcal{P}$, then $gH_xg^{-1} = H_{gx}$.
- *Spinning:* For any $x \neq y \in \mathcal{P}$ and any non-trivial $h \in H_y$, we have $\mathsf{d}^{\pi}_y(x, hx) > L$.

Notice that G acts on $\mathcal{P}$ by isometries, as the projections $\{\pi_Y\}_{Y \in \mathcal{Y}}$ and $\{\pi_x\}_{x \in X}$ are G-equivariant. The following lemma is immediate.

Lemma 3.4. *If $(X, \mathcal{Y}, G, \{H_Y\}_{Y \in \mathcal{Y}})$ satisfies Hypothesis 3.1 with respect to (δ, K, M_0, R, L), then the family of subgroups $\{H_x\}_{x \in \widehat{X}}$ forms an equivariant L–spinning family.*

We gather here some facts about spinning families from [CMM21, CM22]. First, the main theorem of [CMM21] states that the subgroup $N := \langle\langle H_x \rangle\rangle_{x \in \widehat{X}} = \langle\langle H_Y \rangle\rangle_{Y \in \mathcal{Y}}$ normally generated by the spinning family is a (generally infinite) free product:

Theorem 3.5. *Let $(X, \mathcal{Y}, G, \{H_Y\}_{Y \in \mathcal{Y}})$ be as in Hypothesis 3.1, and let $\mathcal{O}$ be any set of orbit representatives for the action of N on $\mathcal{Y}$. Then $N \cong \ast_{Z \in \mathcal{O}} H_Z$.*

We can use Theorem 3.5 to characterize stabilizers in N:

Corollary 3.6. *Let $(X, \mathcal{Y}, G, \{H_Y\}_{Y \in \mathcal{Y}})$ be as in Hypothesis 3.1. For $Y \in \mathcal{Y}$ we have that $N \cap \mathrm{Stab}_G(Y) = H_Y$.*

Proof. Let $\mathcal{O}$ be a collection of N-orbit representatives including Y. Since H_Y is a non-trivial free factor of $N \cong \ast_{Z \in \mathcal{O}} H_Z$, it is malnormal in N, i.e., it intersects any of its conjugates trivially. Furthermore, if $n \in N$ fixes Y, then it normalizes H_Y by Hypothesis 3.1.(6), and so $n \in H_Y$ by malnormality, as required. ☐

Remark 3.7 (Complexity)**.** There is a partial order $\prec$ on elements of N, called *complexity*, which is invariant under conjugation by elements of N. The only facts we will need about this partial order is that it has the identity as its unique minimal element, and that descending chains have finite length; this will allow for inductive arguments on the complexity of an element. We omit the full definition and refer the reader to [CM22, Section 3] for details.

Proposition 3.8 ([CM22, Lemma 3.2])**.** *Assume that $(X, \mathcal{Y}, G, \{H_Y\}_{Y\in\mathcal{Y}})$ satisfies Hypothesis 3.1. Let $x \in \widehat{X}$ and $h \in N$ be such that $hx \neq x$. Then there exist $Y \in \mathcal{Y}$ and $h_Y \in H_Y$ such that the following hold.*

1. *Either $v_Y \in \{x, hx\}$ or $\mathsf{d}^{\pi}_Y(x, hx) > L/10$; and*
2. *$h_Y h \prec h$.*

We say that (Y, h_Y) is a shortening pair *for (x, h).*

Clay and Mangahas further investigated the properties of spinning families and showed that $\mathcal{P}/N$ is hyperbolic, provided that $L > L_{hyp}$ [CM22, Theorem 1.1]. In Section 3.2, we use these properties to prove that $\widehat{X}/N$, rather than the quotient of $\mathcal{P}$, is uniformly hyperbolic, regardless of the choice of L (as long as L is bigger than the constant $\overline{L}$ from Equation (3)).

We conclude the subsection with the following Corollary, which is crucial in the application to quasi-isometric rigidity of random quotients of mapping class groups [Man23]:

Corollary 3.9 (Large injectivity radius)**.** *Under the assumptions of Proposition 3.8 let $\tau = (L/10 - 2(B + JR))/J \geqslant 2$. For every $x \in X$ and every $h \in N - \{1\}$, we have $\mathsf{d}_X(x, hx) > \tau$. In particular, N acts freely on X.*

Proof. Following Definition 2.21, let $\mathcal{V}(x) = \mathcal{V}_R(x)$ be the set of nearest subspaces of $\mathcal{Y}$. This set is non-empty as the action is R-cobounded, so let $Y \in \mathcal{V}(x)$. If $hY = Y$, then $h \in H_Y$ by Corollary 3.6. Hence $\mathsf{d}_Y(x, hx) > L > J$ by Hypothesis 3.1, and so $\mathsf{d}_X(x, hx) > L/J$ by Lemma 2.3.

Suppose instead that $hY \neq Y$. Then by Proposition 3.8 there exists $U \in \mathcal{Y}$ such that $\mathsf{d}_U(Y, hY) > L/10$. Since projections are J-Lipschitz and x is R-close to Y, we have that $\mathsf{d}_U(x, hx) > L/10 - 2(B + JR) > J$, since L is greater than the constant $\overline{L}$ from Equation (3). Hence again Lemma 2.3 yields that $\mathsf{d}_X(x, hx) > (L/10 - 2(B + JR))/J$. ☐

3.2 Hyperbolicity of the quotient graph

We now turn our attention to the quotient $\overline{X} \coloneqq \widehat{X}/N$. Let $q\colon \widehat{X} \to \overline{X}$ be the quotient map. We say that a subgraph $T \subseteq \widehat{X}$ *lifts* a subgraph $\overline{T} \subseteq \overline{X}$ if q restricts to an isometry between T and $\overline{T}$. To avoid confusion, we will denote points in $\overline{X}$ by $\overline{x}$ and points in $\widehat{X}$ (and X) simply by x. By an abuse of notation, we consider points in $\overline{X}$ to be equivalence classes of points in $\widehat{X}$: if $q(x) = \overline{x}$, we will write $x \in \overline{x}$ and say x is a *representative* of $\overline{x}$. We first define a certain class of representatives.

Definition 3.10. Given $\overline{x}, \overline{y} \in \overline{X}$, two points $x \in \overline{x}$ and $y \in \overline{y}$ are *minimal distance representatives*, or simply *minimal*, if

$$\mathsf{d}_{\widehat{X}}(x, y) = \inf_{x' \in \overline{x},\, y' \in \overline{y}} \mathsf{d}_{\widehat{X}}(x', y').$$

Lemma 3.11. *Let $\overline{x}, \overline{y} \in \overline{X}$, and let $x \in \overline{x}$, $y \in \overline{y}$ be minimal representatives. Suppose $\alpha = \alpha_1 * \alpha_2$ is a geodesic in $\widehat{X}$ from x to y with endpoints $(\alpha_1)_+ = v_Y = (\alpha_2)_-$ for some $Y \in \mathcal{Y}$. For any $h_Y \in H_Y$, the path $\alpha' = \alpha_1 * h_Y\alpha_2$ is also a geodesic in $\widehat{X}$.*

Proof. Since $h_Y \in \mathrm{Stab}(Y)$, we have $h_Y v_Y = v_Y$. Therefore α' has the same length as α, and so $\mathsf{d}_{\widehat{X}}(x, h_Y y) \leqslant \mathsf{d}_{\widehat{X}}(x, y)$. On the other hand, since x and y are minimal and $h_Y y$ is also a representative of $\overline{y}$, the distance from $h_Y y$ to x is at least the distance from y to x. Together, this implies that $\mathsf{d}_{\widehat{X}}(x, y) = \mathsf{d}_{\widehat{X}}(x, h_Y y)$, and so α' is a geodesic in $\widehat{X}$. □

We say that we *bend the path α along v_Y* to obtain the path $\alpha' = \alpha_1 * h_Y\alpha_2$, and we call α' a *bent path.* Notice that, since the endpoints of α are minimal, the image in the quotient $q(\alpha)$ is a geodesic between $\overline{x}$ and $\overline{y}$, and both α and α' are lifts of $q(\alpha)$.

The leitmotiv of many arguments throughout this paper is that, if one combines Proposition 3.8 with the strong bounded geodesic image Lemma 2.13, then many combinatorial configurations lift from $\overline{X}$ to $\widehat{X}$. We showcase this in the next Proposition, where we prove that geodesic quadrangles admit lifts. The "moreover" part will be relevant in Section 3.3 to ensure that the image of certain WPD elements of G remain WPD elements of G/N.

Proposition 3.12. *Let $\overline{Q} \subseteq \overline{X}$ be a geodesic quadrangle with vertices $\overline{v}_1, \overline{v}_2, \overline{v}_3, \overline{v}_4$. Then there exists a geodesic quadrangle $Q \subseteq \widehat{X}$ which lifts $\overline{Q}$.*

Moreover, if the geodesics $[\overline{v}_1, \overline{v}_2]$ and $[\overline{v}_3, \overline{v}_4]$ of $\overline{Q}$ have lifts $[v_1', v_2']$ and $[v_3', v_4']$, respectively, such that $\sup_{Y \in \mathcal{Y}} \mathsf{d}^{\pi}_Y(v_1', v_2') \leqslant L/40$ and $\sup_{Y \in \mathcal{Y}} \mathsf{d}^{\pi}_Y(v_3', v_4') \leqslant L/40$, then the lifts of $[\overline{v}_1, \overline{v}_2]$ and $[\overline{v}_3, \overline{v}_4]$ contained in Q can be chosen to be N–translates of $[v_1', v_2']$ and $[v_3', v_4']$, respectively.

Proof. Lift each geodesic side of $\overline{Q}$ to a geodesic in $\widehat{X}$. Up to replacing each lift by some N-translate, this produces a concatenation of four geodesics $\gamma_1 = [v_1, v_2]$, $\gamma_2 = [v_2, v_3]$, $\gamma_3 = [v_3, v_4]$, and $\gamma_4 = [v_4, v_5]$, where $v_5 = hv_1$ for some $h \in N$. Under the hypothesis of the "moreover" part, we can choose γ_1 and γ_3 to be N–translates of $[v_1', v_2']$ and $[v_3', v_4']$, respectively. We call this configuration an *open lift* of $\overline{T}$.

Recall that, by Remark 3.7, descending chains in $\prec$ have finite length, so we proceed by induction on the complexity of h. If $h = 1$, or more generally if $hv_1 = v_1$, then the union $\gamma_1 \cup \ldots \cup \gamma_4$ is already a geodesic quadrangle Q lifting $\overline{Q}$, as required. Otherwise, we show that we can find another open lift of $\overline{Q}$ where the new endpoints differ by some $h' \prec h$, reducing the complexity. Indeed, since $hv_1 \neq v_1$, Proposition 3.8 provides a shortening pair (Y, h_Y) for (v_1, h). There are two cases to consider.

- Suppose first that $v_Y = v_j$ for some $1 \leqslant j \leqslant 5$. For every $i \geqslant j$, replace γ_i by $h_Y\gamma_i$. This results in another open lift of $\overline{T}$, made of N-translates of the original γ_is, but now v_1 and $h_Y hv_1$ differ by $h_Y h$, which has lower complexity than h by definition of the shortening pair.

- Now suppose that $v_Y \neq v_j$ for all j, so that all projections from v_j to Y are defined. By Proposition 3.8 we have that $\mathsf{d}^{\pi}_Y(v_1, hv_1) > L/10$, and the triangle inequality yields that at least one of $\mathsf{d}^{\pi}_Y(v_1, v_2)$, $\mathsf{d}^{\pi}_Y(v_2, v_3)$, $\mathsf{d}^{\pi}_Y(v_3, v_4)$, and $\mathsf{d}^{\pi}_Y(v_4, hv_1)$ is larger than $L/40$. We assume that $\mathsf{d}^{\pi}_Y(v_4, hv_1) > L/40$; an analogous argument holds in the other cases. Since we chose L greater than the constant $\overline{L}$ from Equation (3), the quantity $L/40$ is greater than the constant C from the bounded geodesic image Lemma 2.13. It follows that v_Y lies on the geodesic γ_4. Bend γ_4 at v_Y by h_Y; in other words, apply h_Y to every vertex of the open lift between v_Y and hv_1. See Figure 2. Since the bent path is still a geodesic lift of $[\overline{v}_4, \overline{v}_1]$, this operation produces a new open lift of $\overline{Q}$. Moreover, as before, v_1 and $h_Y h v_1$ now differ by $h_Y h \prec h$. Notice that, in the setting of the "moreover" part, both $\mathsf{d}^{\pi}_Y(v_1, v_2)$ and $\mathsf{d}^{\pi}_Y(v_3, v_4)$ are at most $L/40$, hence the bending procedure replaces each of γ_1 and γ_3 by an N–translate.

In both cases, we conclude by induction: after finitely many steps, we have obtained a geodesic quadrangle Q that lifts $\overline{Q}$, and, in the "moreover" setting, the lifts of $[\overline{v}_1, \overline{v}_2]$ and $[\overline{v}_3, \overline{v}_4]$ contained in Q are N-translates of $[v_1', v_2']$ and $[v_3', v_4']$, respectively. □

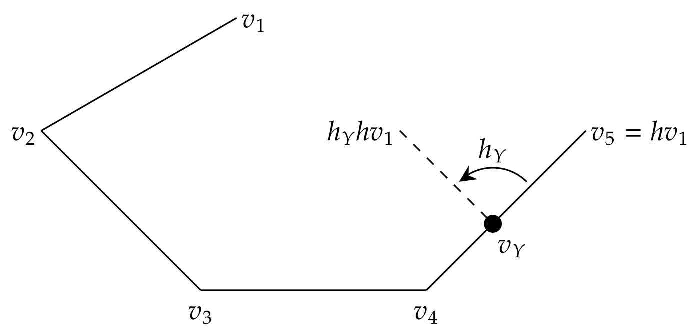

Figure 2: The open lift from the proof of Proposition 3.12, and its bending at the shortening vertex (here, the dashed line).

Theorem 3.13. *If $(X, \mathcal{Y}, G, \{H_Y\}_{Y\in\mathcal{Y}})$ satisfies Hypothesis 3.1, then $\overline{X}$ is $\widehat{\delta}$-hyperbolic, where $\widehat{\delta} = \widehat{\delta}(\delta, K)$ is the hyperbolicity constant of $\widehat{X}$ from Lemma 2.8.*

Proof. Let $\overline{T} \subseteq \overline{X}$ be a geodesic triangle, which we see as a degenerate quadrangle, and let $T \subseteq \widehat{X}$ be a lift of $\overline{T}$. Given any point $\overline{p} \in \overline{T}$ let $p \in T$ be its lift. By slimness of triangles in $\widehat{X}$, there exists $q \in T$ on one of the other sides such that $\mathsf{d}_{\widehat{X}}(p, q) \leqslant \widehat{\delta}$. But then, since the projection map $q\colon \widehat{X} \to \overline{X}$ is 1-Lipschitz, we have that $\mathsf{d}_{\overline{X}}(\overline{p}, \overline{q}) \leqslant \mathsf{d}_{\widehat{X}}(p, q) \leqslant \widehat{\delta}$, where $\overline{q}$ is the projection of q. This proves that $\overline{T}$ is $\widehat{\delta}$-slim, as required. □

3.3 Preserving acylindrical hyperbolicity

In this subsection, we show that quotients of acylindrically hyperbolic groups by spinning families are acylindrically hyperbolic. We go about this by showing that if $G \curvearrowright X$ admits independent loxodromic WPD elements whose axes are "transverse" to the collection $\mathcal{Y}$, then the quotient G/N does, as well.

Recall that, given an isometric action of a group G on a hyperbolic space X, an element $g \in G$ is *loxodromic* if for every $x \in X$ the map $\mathbb{Z} \to X$ mapping n to $g^n(x)$ is a quasi-isometric embedding. In this case, the *limit set* of $\langle g \rangle$ consists of the two endpoints of the quasigeodesic $\{g^n(x)\}_{n\in\mathbb{Z}}$ inside ∂X, which do not depend on x. Two loxodromic isometries of a hyperbolic space are *independent* if their limit sets are disjoint. Furthermore, following [BF02], a loxodromic element $g \in G$ is *weakly properly discontinuous*, or *WPD*, if for every $\varepsilon > 0$ and every $x \in X$ there exists $M \in \mathbb{N}$ such that

$$\left|\left\{h \in G \mid \mathsf{d}_X(x, hx) \leqslant \varepsilon,\, \mathsf{d}_X\left(g^M x, hg^M x\right) \leqslant \varepsilon\right\}\right| < \infty.$$

The core of the arguments in this subsection is the following lemma.

Lemma 3.14. *Let $x, y \in X$ be such that $\sup_{Y \in \mathcal{Y}} \mathsf{d}_Y^{\pi}(x, y) \leqslant L/20$. Then*

$$\frac{1}{L/20 + 2C}\mathsf{d}_X(x, y) \leqslant \mathsf{d}_{\widehat{X}}(x, y) = \mathsf{d}_{\overline{X}}(\overline{x}, \overline{y}). \tag{4}$$

Proof. Let γ be an $\widehat{X}$-geodesic connecting x to y, and let $\widetilde{\gamma}$ be its de-electrification in X. Clearly the length of $\widetilde{\gamma}$ is at least $\mathsf{d}_X(x, y)$; hence the inequality on the left of (4) follows if we bound the length of $\widetilde{\gamma}$ by $(L/20 + 2C)$-times the length of γ. In turn, by the definition of the de-electrification, it suffices to show that whenever γ has a two-edge subsegment of the form $\{c, v_Y, d\}$, where $Y \in \mathcal{Y}$ and $c, d \in Y$, then $\mathsf{d}_X(c, d) \leqslant L/20 + 2C$. To see this, assume that c is closer to x than d. Since $\gamma|_{[x,c]}$ does not pass through v_Y, $\mathsf{d}_Y^{\pi}(c, x) \leqslant C$ by Lemma 2.13, and symmetrically $\mathsf{d}_Y^{\pi}(d, y) \leqslant C$. Therefore, since $c, d \in Y$, we have that

$$\mathsf{d}_X(c, d) = \mathsf{d}_Y^{\pi}(c, d) \leqslant \mathsf{d}_Y^{\pi}(c, x) + \mathsf{d}_Y^{\pi}(x, y) + \mathsf{d}_Y^{\pi}(y, d) \leqslant L/20 + 2C.$$

For the equality on the right of (4), suppose x and some $y' \in \overline{y}$ are minimal and let η be an $\widehat{X}$-geodesic connecting x to y'. We claim that γ and η have the same length. To see this, let $h \in N$ be such that $hy = y'$. We proceed by induction on the complexity of h. If h fixes y we are done. Otherwise, let (Y, h_Y) be a shortening pair for (y, h), as in Proposition 3.8. We have that

$$L/10 < \mathsf{d}_Y^{\pi}(y, y') \leqslant \mathsf{d}_Y^{\pi}(y, x) + \mathsf{d}_Y^{\pi}(x, y') \leqslant L/20 + \mathsf{d}_Y^{\pi}(x, y').$$

Since $L/10 - L/20 = L/20 > C$, Lemma 2.13 implies that v_Y lies on η. Bending η at v_Y, we conclude by induction. □

We are now ready to prove the main technical result of this subsection.

Proposition 3.15. *Let $f, g \in G$ be independent loxodromic isometries for the action on X, and let $x \in X$ be such that*

$$\sup_{Y \in \mathcal{Y}} \sup_{m,n \in \mathbb{Z}} \mathsf{d}_Y^{\pi}(f^m x, g^n x) \leqslant L/40.$$

Then $\overline{f}, \overline{g} \in G/N$ are independent loxodromics for the action on $\overline{X}$. Moreover, if f is a WPD element, then so is $\overline{f}$.

Proof. The hypotheses, together with Lemma 3.14, imply that $\mathsf{d}_{\overline{X}}(\overline{x}, \overline{f}^n \overline{x})$ and $\mathsf{d}_{\overline{X}}(\overline{x}, \overline{g}^n \overline{x})$ both grow linearly in n, so $\overline{f}$ and $\overline{g}$ are loxodromic. Moreover, the same Lemma 3.14 yields that

$$\mathsf{d}_{Haus}(\langle \overline{f} \rangle \overline{x}, \langle \overline{g} \rangle \overline{x}) \geqslant \frac{1}{L/20 + 2C}\mathsf{d}_{Haus}(\langle f \rangle x, \langle g \rangle x) = \infty,$$

so $\overline{f}$ and $\overline{g}$ are indeed independent.

Now suppose that f is a WPD element with respect to the action of G on X. By hypothesis, a quasiaxis for f can fellow travel along a subspace $Y \in \mathcal{Y}$ only for a uniformly bounded amount, so [MMS24, Theorem 2.4] implies that f is a WPD element with respect to the action of G on $\widehat{X}$. Now fix $\varepsilon > 0$ and $\overline{x} \in \overline{X}$, let $x \in \widehat{X}$ be a lift of $\overline{x}$, and let $M > 0$ be such that the set

$$\Delta = \{h \in G \mid \mathsf{d}_{\widehat{X}}(x, hx) \leqslant \varepsilon,\ \mathsf{d}_{\widehat{X}}(f^M x, fg^M x) \leqslant \varepsilon\}$$

is finite. Let D denote the size of Δ.

Suppose there exist distinct elements $\overline{g}_1, \ldots, \overline{g}_p \in G/N$ such that

$$\mathsf{d}_{\overline{X}}(\overline{x}, \overline{g}_i\overline{x}) \leqslant \varepsilon \quad \text{and} \quad \mathsf{d}_{\overline{X}}(\overline{f}^M\overline{x}, \overline{g}_i\overline{f}^M\overline{x}) \leqslant \varepsilon.$$

Consider the geodesic quadrilateral $\overline{Q}$ in $\overline{X}$ with vertices $\overline{x}, \overline{f}^M\overline{x}, \overline{g}_i\overline{f}^M\overline{x}$, and $\overline{f}^M\overline{x}$. By the "moreover" statement of Proposition 3.12, there exist $g_i \in \overline{g}_i$ such that the quadrilateral $\overline{Q}$ lifts to a geodesic quadrilateral Q of $\widehat{X}$ with vertices x, f^Mx, g_if^Mx, and g_ix. In particular, we have

$$\mathsf{d}_{\widehat{X}}(x, g_ix) = \mathsf{d}_{\overline{X}}(\overline{x}, \overline{g}_i\overline{x}) \leqslant \varepsilon \quad \text{and, similarly,} \quad \mathsf{d}_{\widehat{X}}(f^Mx, g_if^Mx) \leqslant \varepsilon.$$

It follows that $g_i \in \Delta$ for each $1 \leqslant i \leqslant p$, and hence $p < D$. This shows that the element $\overline{f}$ is WPD with respect to the action of G/N on $\overline{X}$, as required. ☐

Corollary 3.16. *Let $(X, \mathcal{Y}, G, \{H_Y\}_{Y\in\mathcal{Y}})$ satisfy Hypothesis 3.1. Suppose there exist independent loxodromics $f, g \in G$ as in Proposition 3.15, and that f is a WPD element. Then G/N is acylindrically hyperbolic.*

Proof. By Proposition 3.15, G/N acts with independent loxodromic isometries on $\overline{X}$ and so is not virtually cyclic. Furthermore, $\overline{f}$ is a loxodromic WPD element for this action. Therefore G/N is acylindrically hyperbolic by [Osi16, Theorem 1.2]. ☐

4 Hierarchical hyperbolicity of spinning quotients

Up until this point, we have worked with an arbitrary hyperbolic graph X with a group action, along with an equivariant collection of uniformly quasiconvex subspaces. In this section, we specialize to the case when X arises as the top level hyperbolic space in a (relative) hierarchically hyperbolic group structure.

4.1 Background on hierarchical hyperbolicity

This subsection gathers definitions and properties of (relative) hierarchically hyperbolic spaces and groups introduced in [BHS19]; see [Rus20, Definition 2.8] for this formulation of some of the axioms.

Definition 4.1 (Relative hierarchically hyperbolic space)**.** Let $E > 0$, and let $\mathcal{X}$ be an (E, E)–quasigeodesic space. A *relatively hierarchically hyperbolic space (relative HHS) structure with constant* E for $\mathcal{X}$ is an index set $\mathfrak{S}$ and a set $\{\mathcal{C}W \mid W \in \mathfrak{S}\}$ of geodesic spaces $(\mathcal{C}W, \mathsf{d}_W)$ such that the following axioms are satisfied.

(1) **(Projections.)** For each $W \in \mathfrak{S}$, there exists a *projection* $\pi_W \colon \mathcal{X} \to 2^{\mathcal{C}W}$ that is a (E, E)–coarsely Lipschitz, E–coarsely onto, E–coarse map.

(2) **(Nesting.)** If $\mathfrak{S} \neq \varnothing$, then $\mathfrak{S}$ is equipped with a partial order $\sqsubseteq$ and contains a unique $\sqsubseteq$–maximal element. When $V \sqsubseteq W$, we say V is *nested* in W. For each $W \in \mathfrak{S}$, we denote by $\mathfrak{S}_W$ the set of all $V \in \mathfrak{S}$ with $V \sqsubseteq W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \sqsubsetneq W$ there is a specified non-empty subset $\rho^V_W \subseteq \mathcal{C}W$ with $\mathsf{diam}(\rho^V_W) \leqslant E$.

(3) **(Finite complexity.)** Any set of pairwise $\sqsubseteq$–comparable elements has cardinality at most E.

(4) **(Orthogonality.)** The set $\mathfrak{S}$ has a symmetric relation called *orthogonality*. If V and W are orthogonal, we write $V \perp W$ and require that V and W are not $\sqsubseteq$–comparable. Further, whenever $V \sqsubseteq W$ and $W \perp U$, we require that $V \perp U$. We denote by $\mathfrak{S}_W^{\perp}$ the set of all $V \in \mathfrak{S}$ with $V \perp W$.

(5) **(Containers.)** For each $W \in \mathfrak{S}$ and $U \in \mathfrak{S}_W$ with $\mathfrak{S}_W \cap \mathfrak{S}_U^{\perp} \neq \varnothing$, there exists $Q \in \mathfrak{S}_W$ such that $V \sqsubseteq Q$ whenever $V \in \mathfrak{S}_W \cap \mathfrak{S}_U^{\perp}$. We call Q the *container for U inside W*.

(6) **(Transversality.)** If $V, W \in \mathfrak{S}$ are neither orthogonal nor $\sqsubseteq$-comparable, we say V and W are *transverse*, denoted $V \pitchfork W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \pitchfork W$, there are non-empty sets $\rho_W^V \subseteq \mathcal{C}W$ and $\rho_V^W \subseteq \mathcal{C}V$, each of diameter at most E.

(7) **(Consistency.)** For all $x \in \mathcal{X}$ and $V, W, U \in \mathfrak{S}$:

- if $V \pitchfork W$, then $\min\left\{\mathsf{d}_W(\pi_W(x), \rho_W^V), \mathsf{d}_V(\pi_V(x), \rho_V^W)\right\} \leqslant E$, and
- if $U \sqsubseteq V$ and either $V \sqsubsetneq W$ or $V \pitchfork W$ and $W \not\perp U$, then $\mathsf{d}_W(\rho_W^U, \rho_W^V) \leqslant E$.

(8) **(Hyperbolicity)** For each $W \in \mathfrak{S}$, either W is $\sqsubseteq$–minimal or $\mathcal{C}W$ is E–hyperbolic.

(9) **(Bounded geodesic image.)** For all $V, W \in \mathfrak{S}$ and for all $x, y \in \mathcal{X}$, if $V \sqsubsetneq W$ and $\mathsf{d}_V(\pi_V(x), \pi_V(y)) \geqslant E$, then every $\mathcal{C}W$–geodesic from $\pi_W(x)$ to $\pi_W(y)$ must intersect $\mathcal{N}_E(\rho_W^V)$.

(10) **(Partial realization.)** If $\{V_i\}$ is a finite collection of pairwise orthogonal elements of $\mathfrak{S}$ and $p_i \in \mathcal{C}V_i$ for each i, then there exists $x \in \mathcal{X}$ so that:

- $\mathsf{d}_{V_i}(\pi_{V_i}(x), p_i) \leqslant E$ for all i;
- for each i and each $W \in \mathfrak{S}$, if $V_i \sqsubsetneq W$ or $W \pitchfork V_i$, we have $\mathsf{d}_W(\pi_W(x), \rho_W^{V_i}) \leqslant E$.

(11) **(Uniqueness.)** There is a function $\theta\colon [0,\infty) \to [0,\infty)$ so that for all $r \geqslant 0$, if $x, y \in \mathcal{X}$ and $\mathsf{d}_{\mathcal{X}}(x,y) \geqslant \theta(r)$, then there exists $W \in \mathfrak{S}$ such that $\mathsf{d}_W(\pi_W(x), \pi_W(y)) \geqslant r$.

(12) **(Large links.)** For all $W \in \mathfrak{S}$ and $x, y \in \mathcal{X}$, there exists $\{V_1, \ldots, V_m\} \subseteq \mathfrak{S}_W - \{W\}$ such that m is at most $E\mathsf{d}_W(\pi_W(x), \pi_W(y)) + E$, and for all $U \in \mathfrak{S}_W - \{W\}$, either $U \in \mathfrak{S}_{V_i}$ for some i, or $\mathsf{d}_U(\pi_U(x), \pi_U(y)) \leqslant E$.

We use $\mathfrak{S}$ to denote the relative HHS structure, including the index set $\mathfrak{S}$, spaces $\{\mathcal{C}W : W \in \mathfrak{S}\}$, projections $\{\pi_W : W \in \mathfrak{S}\}$, and relations $\sqsubseteq$, $\perp$, $\pitchfork$. We call an element $U \in \mathfrak{S}$ a *domain*, the associated space $\mathcal{C}U$ its *coordinate space*, and call the maps ρ_W^V the *relative projections* from V to W. The number E is called the *hierarchy constant* for $\mathfrak{S}$; notice that every $E' \geqslant E$ is again a hierarchy constant for $\mathfrak{S}$, so we are often free to enlarge E by a bounded amount.

A relative HHS is called a *hierarchically hyperbolic space* (HHS) if for every $W \in \mathfrak{S}$ the space $\mathcal{C}W$ is E–hyperbolic, in this case we say $\mathfrak{S}$ is a *HHS structure* on $\mathcal{X}$.

A quasigeodesic space $\mathcal{X}$ is a *(relative) HHS with constant E* if there exists a (relative) HHS structure on $\mathcal{X}$ with constant E. The pair $(\mathcal{X}, \mathfrak{S})$ denotes a (relatively) HHS equipped with the specific HHS structure $\mathfrak{S}$.

The large links axiom (Definition 4.1.(12)) can be replaced with the following, which is traditionally called *passing up*.

(12') **(Passing up).** For every $t > 0$, there exists an integer $P = P(t) > 0$ such that if $V \in \mathfrak{S}$ and $x, y \in \mathcal{X}$ satisfy $\mathsf{d}_{U_i}(x, y) > E$ for a collection of domains $\{U_i\}_{i=1}^{P}$ with $U_i \in \mathfrak{S}_V$, then there exists $W \in \mathfrak{S}_V$ with $U_i \subsetneq W$ for some i such that $\mathsf{d}_W(x, y) > t$.

It was shown in [BHS19, Lemma 2.5] that every HHS satisfies the passing up axiom. The following lemma, which is the converse implication, is stated explicitly in [Dur23, Section 4.8], but the strategy behind the proof appears in [PS23, Lemma 5.3]. The proof is written for a HHS, but it does not use the hyperbolicity of the spaces $\mathcal{C}W$, and so it also applies to a relative HHS.

Lemma 4.2. *If $(\mathcal{X}, \mathfrak{S})$ satisfies axioms (1)–(11) from Definition 4.1, as well as the passing up axiom (12'), then $(\mathcal{X}, \mathfrak{S})$ is a relative HHS. If moreover the spaces $\mathcal{C}W$ are E–hyperbolic for all $W \in \mathfrak{S}$, then $(\mathcal{X}, \mathfrak{S})$ is a HHS.*

The following is a combination of [BHS17a, Lemma 1.8] and the consistency axiom (7).

Lemma 4.3. *Let $(\mathcal{X}, \mathfrak{S})$ be a relative HHS. If $U, V \in \mathfrak{S}_W$ are not transverse, then $\mathsf{d}_S(\rho_W^U, \rho_W^V)$ is at most $2E$.*

A hallmark of hierarchically hyperbolic spaces is that every pair of points can be joined by a special family of quasigeodesics called *hierarchy paths*, each of which projects to a quasigeodesic in each of the spaces $\mathcal{C}W$.

Definition 4.4. A *λ–hierarchy path* γ in a HHS $(\mathcal{X}, \mathfrak{S})$ is a λ–quasigeodesic with the property that $\pi_W \circ \gamma$ is an unparametrized λ–quasigeodesic for each $W \in \mathfrak{S}$.

Theorem 4.5 ([BHS19, Theorem 6.11])**.** *Let $(\mathcal{X}, \mathfrak{S})$ be a relative HHS with constant E. There exist $\lambda \geqslant 1$ depending only on E so that every pair of points in $\mathcal{X}$ is joined by a λ–hierarchy path.*

Given $A, B, C, D \in \mathbb{R}$, write $A \preceq_{C,D} B$ to mean that $A \leqslant BC + D$, and $A \asymp_{C,D} B$ if $B \preceq_{C,D} A \preceq_{C,D} B$. Let $\{\!\{A\}\!\}_B$ be the quantity which is A if $A \geqslant B$, and is 0 otherwise.

Theorem 4.6 (Distance formula, [BHS19, Theorem 6.10])**.** *Let $(\mathcal{X}, \mathfrak{S})$ be a relative HHS. There exists s_0 such that for all $s \geqslant s_0$, there exist k_1, k_2 so that for all $x, y \in \mathcal{X}$,*

$$\mathsf{d}_{\mathcal{X}}(x, y) \asymp_{k_1, k_2} \sum_{U \in \mathfrak{S}} \{\!\{\mathsf{d}_U(x, y)\}\!\}_s .$$

We next introduce the notion of a hierarchically hyperbolic group.

Definition 4.7 (Hierarchically hyperbolic group)**.** Let G be a finitely generated group and $\mathcal{X}$ be the Cayley graph of G with respect to some finite generating set. We say G is a (relatively) *hierarchically hyperbolic group* (HHG) if the following hold.

(i) The space $\mathcal{X}$ admits a (relative) HHS structure $\mathfrak{S}$ with hierarchy constant E.

(ii) There is a $\sqsubseteq$–, $\perp$–, and $\pitchfork$–preserving action of G on $\mathfrak{S}$ by bijections such that $\mathfrak{S}$ contains finitely many G–orbits.

(iii) For each $W \in \mathfrak{S}$ and $g \in G$, there exists an isometry $g_W : \mathcal{C}W \to \mathcal{C}(gW)$ satisfying the following for all $V, W \in \mathfrak{S}$ and $g, h \in G$.

- The map $(gh)_W : \mathcal{C}W \to \mathcal{C}(ghW)$ is equal to the map $g_{hW} \circ h_W : \mathcal{C}W \to \mathcal{C}(hW)$.
- For each $x \in X$, $g_W(\pi_W(x))$ and $\pi_{gW}(g \cdot x)$ are at distance at most E in $\mathcal{C}(gW)$.
- If $V \pitchfork W$ or $V \subsetneq W$, then $g_W(\rho^V_W)$ and ρ^{gV}_{gW} are at distance at most E in $\mathcal{C}(gW)$.

The structure $\mathfrak{S}$ satisfying (i)—(iii) is called a (relative) hierarchically hyperbolic group structure on G. We use $(G, \mathfrak{S})$ to denote G equipped with a specific (relative) HHG structure on G. In Item (iii) we often drop the subscript on the isometry g_W and simply write g when the domain W is clear from context.

Remark 4.8. The second and third bullet points in Item (iii) imply that the isometry g_W is coarsely equivariant with respect to projections π_W and the ρ–maps. In fact, we can assume that these isometries are genuinely equivariant; see [DHS20, Section 2.1]. That is, we assume that for all $W \in \mathfrak{S}$, all $g \in G$, and all $V \in \mathfrak{S}$ with $V \pitchfork W$ or $V \subsetneq W$, we have:

- $g_W(\pi_W(x)) = \pi_{gW}(g \cdot x)$, and
- $g_W(\rho^V_W) = \rho^{gV}_{gW}$.

Remark 4.9 (Convention on $\mathcal{C}S$)**.** In a HHS structure, one can assume that $\mathcal{C}W$ is a graph for every $W \in \mathfrak{S}$, as it can always be $\mathrm{Stab}_G(W)$-equivariantly replaced by a graph (see, e.g., [CdlH16, Lemma 3.B.6]). Furthermore, if G is a HHG, the projection $\pi_S : G \to \mathcal{C}S$ can be assumed to be a bijection. This is because, since π_S is E-coarsely onto and G-equivariant, G acts coboundedly on $\mathcal{C}S$, so $\mathcal{C}S$ is G-equivariantly quasi-isometric to a Cayley graph of G with respect to some possibly infinite generating set. Hence we can (and will) identify points of $\mathcal{C}S$ with elements of G. Finally, if $\mathcal{C}S$ is bounded, we will always assume that the HHG constant E is larger than $\mathsf{diam}(\mathcal{C}S)$.

For the next lemma, recall that the action of a group G on a metric space X is *acylindrical* if for all $\varepsilon > 0$ there exist constants $R = R(\varepsilon) \geqslant 0$ and $N = N(\varepsilon) \geqslant 0$ such that for every $x, y \in X$ with $\mathsf{d}(x, y) \geqslant R$, we have

$$\#\{g \in G \mid \mathsf{d}(x, gx) \leqslant \varepsilon \text{ and } \mathsf{d}(y, gy) \leqslant \varepsilon\} \leqslant N.$$

A group is *acylindrically hyperbolic* if it admits a non-elementary acylindrical action on a hyperbolic space, that is, an acylindrical action that contains two independent loxodromic isometries.

Lemma 4.10. *If $(G, \mathfrak{S})$ is a (relative) HHG whose top-level coordinate space $X \coloneqq \mathcal{C}S$ is hyperbolic, then G acts acylindrically on X. As a consequence, if G is not virtually cyclic and X is unbounded, then G is acylindrically hyperbolic, and we say $(G, \mathfrak{S})$ is an* acylindrically hyperbolic (relative) HHG*.*

Proof. This is [BHS17b, Corollary 14.4], which is stated for HHGs but whose proof does not use the hyperbolicity of non-maximal elements of $\mathfrak{S}$. □

The following lemma follows immediately from the distance formula; see also [ABD21] in the case of a HHG.

Lemma 4.11. *Let $(G, \mathfrak{S})$ be a relative HHG and suppose $H \leqslant G$ is (λ, c)–quasi-isometrically embedded by the orbit map $G \to \mathcal{C}S$. There is a constant $\aleph = \aleph(\mathfrak{S}, \lambda, c)$ such that the diameter of $\pi_U(H)$ is at most $\aleph$ for all $U \in \mathfrak{S} - \{S\}$.*

4.2 Statement of the main result

Given a metric space X and a group G, we say that an isometric action $G \curvearrowright X$ is *geometric* if it is cobounded and *metrically proper*, i.e., for every $x \in X$ and every $R > 0$ the set $\{g \in G \mid \mathsf{d}(x, gx) \leqslant R\}$ is finite. For the rest of the section, we shall work under the following strengthening of Hypothesis 3.1:

Hypothesis 4.12. Let G be a (relative) HHG, with top-level coordinate space X and HHG constant $E > 0$. Assume that $(X, \mathcal{Y}, G, \{H_Y\}_{Y \in \mathcal{Y}})$ satisfy Hypothesis 3.1 with respect to (E, K, M_0, R, L). Notice that we can always choose $R = 0$ by Remark 4.9. Furthermore, suppose that:

7. G acts cofinitely on $\mathcal{Y}$;
8. For each $Y \in \mathcal{Y}$, the subgroup H_Y acts geometrically on Y; and
9. $L > \widetilde{L}$, where $\widetilde{L} = \widetilde{L}(E, K, M_0)$ is defined in Equation (6) below and is bounded linearly in M_0.

We say that the collection $(X, \mathcal{Y}, G, \{H_Y\}_{Y \in \mathcal{Y}})$ satisfies Hypothesis 4.12 with respect to (E, K, M_0, L). We drop the constants when unimportant or clear from context. Let $N \coloneqq \langle\!\langle H_Y \rangle\!\rangle_{Y \in \mathcal{Y}}$ be the normal subgroup generated by the spinning family.

Our main result is a technical formulation of Theorem E that quotients of (relative) HHGs by subgroups forming a sufficiently spinning family are (relative) HHGs.

Theorem 4.13. *If G is a (relative) HHG such that $(X, \mathcal{Y}, G, \{H_Y\}_{Y \in \mathcal{Y}})$ satisfies Hypothesis 4.12, then G/N is a (relative) HHG.*

Remark 4.14. Notice that Theorem 4.13 holds trivially when G is *not* an acylindrically hyperbolic (relative) HHG. Indeed, by Lemma 4.10, either such a G is virtually cyclic, or its top-level coordinate space X is bounded. In the first case, any quotient of G is still virtually cyclic, hence hierarchically hyperbolic. In the second case, since $\widetilde{L} \geqslant E$ by Equation (6), and since we are assuming that $E \geqslant \mathsf{diam}(X)$ by Remark 4.9, then the only L-spinning family of subgroups of G is the trivial family, so that $N = \{1\}$ and $G/N = G$. In light of this, the bulk of work is to deal with the case that X is non-elementary hyperbolic, which we shall assume for the remainder of the section.

Assuming that G is acylindrically hyperbolic does not guarantee that the quotient is again acylindrically hyperbolic. For example, if G is a surface group, X is a Cayley graph of G, $Y = X$, and H is a normal subgroup of sufficiently large finite index, then (X, Y, G, H) satisfies Hypothesis 4.12, but the quotient is finite. However, under the additional assumption that there exist two independent loxodromics whose axes have uniformly bounded projections to all $Y \in \mathcal{Y}$, Proposition 3.15 will ensure that G/N is again an acylindrically hyperbolic HHG. We note that this condition is equivalent to G having two independent loxodromic elements for the action on $G \curvearrowright \widehat{X}$, where $\widehat{X}$, defined in the next subsection, is a slight modification of the cone-off from Definition 2.7.

4.3 A modified cone-off

In describing the quotient hierarchy structure, we will use both projections coming from the relative HHS structure and the projection complex structure described in Corollary 2.22. It will be convenient to first modify the space X by a quasi-isometry that introduces new cone points for each domain $U \in \mathfrak{S} - \{S\}$ that will serve as an anchor for canonically defining projections and the quotient action. We first set

$$A \coloneqq \max\{K, 3E\} + 4E + D \tag{5}$$

where $D = D(E, \max\{K, 3E\})$ is defined as in Lemma 2.9. Let X' be formed from X by coning off the collection of subspaces $\{\mathcal{N}_A(\rho_S^U) \mid U \in \mathfrak{S} - \{S\}\}$ such that the cone point over $\mathcal{N}_A(\rho_S^U)$ is v_U. Then G acts acylindrically, 1-coboundedly, and by isometries on X'.

Furthermore, the natural inclusion map $X \to X'$ is a uniform quality quasi-isometry, since the sets we are coning off have diameter which is bounded in terms of E and K; in particular X' is $E'(E, K)$–hyperbolic. For the same reason, each $Y \in \mathcal{Y}$, identified with its image under the inclusion, is $K'(E, K)$–quasiconvex in X', so there is a projection $\pi'_Y : X' \to Y$ which differs from $\pi_Y : X \to Y$ by a uniformly bounded amount, depending on E and K. As a consequence, using that π'_Y is uniformly Lipschitz with respect to E and K, for every $U \in \mathfrak{S}$, the set $\pi'_Y(v_U)$ is within uniformly bounded Hausdorff distance from $\pi_Y(\rho_S^U)$. We thus expand the domain of the projection π_Y to the whole X' by setting

$$\pi_Y(v_U) \coloneqq \pi_Y(\rho_S^U).$$

Notice also that $\mathcal{Y}$ is M_0'–geometrically separated in X', for some M_0' which differs from M_0 by a uniform amount; in particular, M_0' is bounded linearly in M_0. Hence $(X', \mathcal{Y})$ again satisfies Hypothesis 2.19 with respect to the constants (E', K', M_0'). With a little abuse of notation, we once and for all replace E with the maximum of E' and the original E, and similarly for K and M_0, so that $(X', \mathcal{Y})$ satisfies Hypothesis 2.19 with respect to (E, K, M_0). This allows us to define the constants B, J, etc. from Section 2.1, all of which are bounded linearly in M_0' and therefore in M_0.

Let $\widehat{X}$ be formed by coning off $\mathcal{Y}$ in X' (not in X), and let v_Y denote the cone points over $Y \in \mathcal{Y}$. Then $\widehat{X}$ satisfies the strong bounded geodesic image property: applying Lemma 2.13 in X', there is a constant $C' = C'(E, K, M_0)$ which is bounded linearly in M_0, and such that if $x, y \in X'$ satisfy $\mathsf{d}_Y^{\pi'}(x, y) > C'$, where the projection to Y is measured in X', then any $\widehat{X}$ geodesic $[x, y]$ passes through the cone point v_Y. However, as the projections in X' and X differ by a uniform amount (depending on E and K), we immediately obtain the following version of strong bounded geodesic image property, no longer mentioning X' and π'_Y:

Lemma 4.15. *There is a constant $C = C(E, K, M_0)$ which is bounded linearly in M_0, such that for every $Y \in \mathcal{Y}$ and $x, y \in \widehat{X} - \{v_Y\}$, if $\mathsf{d}_Y^{\pi}(x, y) \geqslant C$ (where the projection distance is measured in X), then any $\widehat{X}$–geodesic $[x, y]$ passes through the cone point v_Y.*

With a slight abuse of notation we still call the above constant C, as the one from Lemma 2.13, since we can just take C to be the maximum of the two constants and ensure that both lemmas hold.

Notice that the collection $\{H_Y\}_{Y \in \mathcal{Y}}$ is still an L'–spinning family on X' (with respect to the original projections π_Y), for some constant L' which differs from L by a bounded amount $\text{Ш} = \text{Ш}(E, K)$. In particular, there exists some constant $L_1(E, K, M_0)$, which is bounded linearly in M_0, such that if $L > L_1$ then L' satisfies Hypothesis 3.1, so that all the consequences from Section 3 still hold for X' and $\widehat{X}$ (with respect to the original projections π_Y). Set

$$\widetilde{L} = \max\{L_1, 100C + \text{Ш}, 20(C + EJ) + \text{Ш}\}, \tag{6}$$

which is still bounded linearly in M_0. From now on we assume that $L > \widetilde{L}$, and in particular

$$L' \geqslant L - \text{III} > \max\{100C, 20(C + EJ)\}. \tag{7}$$

We conclude this subsection with some remarks about the construction.

Remark 4.16 (Projection to domains)**.** Recall that, by Remark 4.9, we are assuming that the vertex set of X coincides with G. In particular, for every $x \in X$, the projection $\pi_U(x)$ is well-defined for every $U \in \mathfrak{S}$.

Remark 4.17 (H_Y–orbits have bounded projection)**.** Since each H_Y acts geometrically on the corresponding $Y \subseteq X$, and since G acts on $\mathcal{Y}$ with finitely many orbits, Lemma 4.11 implies the existence of a constant $\aleph$ such that $\mathsf{diam}(\pi_U(H_Y \cdot y)) \leqslant \aleph$ for every $Y \in \mathcal{Y}$, every $y \in Y$, and every $U \in \mathfrak{S} - \{S\}$.

Remark 4.18. For every $U \in \mathfrak{S} - \{S\}$, the set $\mathcal{N}_A(\rho_S^U)$ is $3E$–quasiconvex by Lemma 2.2, regardless of the value of A. Hence $\widehat{X}$ can be seen as the cone-off of X with respect to the family $\mathcal{H} = \{\mathcal{N}_A(\rho_S^U)\}_{U \in \mathfrak{S} - \{S\}} \cup \mathcal{Y}$, whose elements are $\max\{K, 3E\}$-quasiconvex. Hence Lemma 2.9 implies that, for every $x, y \in X$, every X-geodesic $[x, y]$, and every $\widehat{X}$-geodesic γ with the same endpoints, we have $[x, y] \subseteq \mathcal{N}_D^X(\widetilde{\gamma})$, where $\widetilde{\gamma}$ is the de-electrification of γ with respect to $\mathcal{H}$ and $D = D(E, \max\{K, 3E\})$ is as in the definition of A.

4.4 Minimal representatives for points in the quotient

The main goal of this subsection is to define canonical lifts of pairs of points in the quotient and to verify that they are well-behaved with respect to projections; see Proposition 4.24 and Proposition 4.25.

Let $\overline{X} \coloneqq \widehat{X}/N$. The composition $X \overset{i}{\hookrightarrow} \widehat{X} \overset{q}{\to} \overline{X}$, where i is the inclusion map and q is the quotient map, is 1–Lipschitz as both i and q are. For each point $x \in \widehat{X}$, including the case $x = v_Y, v_U \in \widehat{X}$, let $\overline{x} \coloneqq q(x)$. For each $U \in \mathfrak{S}$ we denote its image in $\mathfrak{S}/N$ by $\overline{U}$; as we will see in Section 4.5, $\mathfrak{S}/N$ will be the index set for the (relative) HHG structure on G/N. The following is the analogue of Definition 3.10:

Definition 4.19 (Minimal representatives)**.** Let $x, y \in X$. We say $\{x, y\}$ are *minimal distance representatives*, or simply *minimal*, if

$$\mathsf{d}_{\widehat{X}}(x, y) = \min_{x' \in \overline{x},\, y' \in \overline{y}} \mathsf{d}_{\widehat{X}}(x', y').$$

If in the above definition we replace x by v_U, for some $U \in \mathfrak{S} - \{S\}$, we say that $\{U, y\}$ are minimal. If we also replace y by v_V for some $V \in \mathfrak{S} - \{S\}$, we say that $\{U, V\}$ are minimal.

Our first goal is to show that if $\{x, U\}$ and $\{x', U\}$ are minimal, then there is a uniform bound on $\mathsf{d}_U(x, x')$. We begin with some preparatory lemmas which investigate the uniqueness of certain lifts in $\widehat{X}$.

Lemma 4.20. *Let $x, y \in \widehat{X}$ be adjacent vertices. For every $x' \in \overline{x}$ and $y' \in \overline{y}$ which are adjacent in $\widehat{X}$, there exists $n \in N$ mapping x to x' and y to y'. In other words, an edge of $\overline{X}$ admits a unique N–orbit of lifts.*

Proof. Up to the action of N, we assume that $y = y'$. Let $h \in N$ be such that $x' = hx$, and let γ be the path x, y, x'. Suppose that $hx \neq x$ and let (Y, h_Y) be a shortening pair, as in Proposition 3.8, so that $L'/10 < \mathsf{d}^{\pi}_{Y}(x, x')$. If $v_Y \notin \{x, y, x'\}$ then using the bounded geodesic image Lemma 4.15, we have

$$L'/10 < \mathsf{d}^{\pi}_{Y}(x, x') \leqslant \mathsf{d}^{\pi}_{Y}(x, y) + \mathsf{d}^{\pi}_{Y}(y, x') \leqslant 2C.$$

However this contradicts the lower bound on L' from Equation (7). Hence it must be that $v_Y \in \{x, y, x'\}$. If we apply h_Y to all vertices of γ between v_Y and x', we obtain a new configuration of N-translates of $\{x, y\}$ and $\{y, x'\}$, where the new x and x' differ by $h_Y h \prec h$. We conclude by induction on the complexity of h. □

Lemma 4.21. *Let γ be a concatenation of two $\widehat{X}$-geodesics, and suppose that no point on γ is of the form v_Y for some $Y \in \mathcal{Y}$, except possibly the endpoints x and x'. If $\overline{x} = \overline{x}'$ then $x = x'$.*

In other words, the only way to get distinct minimal representatives in the same orbit is to bend a geodesic at a cone point.

Proof. Let $h \in N$ be such that $x' = hx$. If $hx = x$, for example, when $h = 1$, we are done; otherwise let (Y, h_Y) be a shortening pair, as in Proposition 3.8. If $v_Y \notin \{x, x'\}$, then neither of the geodesics forming γ passes through v_Y, so the diameter of the projection of each geodesic is at most C by Lemma 4.15. Thus $L'/10 \leqslant 2C$, contradicting Equation (7). Therefore we must have that $v_Y \in \{x, x'\}$. Bend the path γ at v_Y. Note that if $v_Y = x$, then we are applying h_Y to all of γ, and x is fixed, while if $v_Y = x'$, then we are only applying h_Y to x', and the whole path γ is fixed. In either case, we have applied an isometry to γ, and so $\mathsf{d}(x, x') = \mathsf{d}(x, h_Y x')$. Since $h_Y h \prec h$, induction shows that we eventually obtain a path with both endpoints at x. Therefore, $\mathsf{d}(x, x')$ must have been zero at every step of the process, so in particular $x' = x$. □

Lemma 4.22. *Let $x, x' \in X$, let γ be an $\widehat{X}$–geodesic between them, and let $V \in \mathfrak{S} - \{S\}$ be such that $\mathsf{d}_V(x, x') > E$. Then $\mathsf{d}_{\widehat{X}}(v_V, \gamma) \leqslant 2$.*

Proof. Since $\mathsf{d}_V(x, x') > E$, the bounded geodesic image in $(G, \mathfrak{S})$ Definition 4.1.(9) implies that, for every X–geodesic α from x to x', there exists $w \in \alpha$ with $\mathsf{d}_X(w, \rho^V_S) \leqslant E$. In turn, by Remark 4.18, α is contained in the D–neighborhood of $\widetilde{\gamma}$ the de-electrification of γ, so we can find $z' \in \widetilde{\gamma}$ such that $\mathsf{d}_X(z', \rho^V_S) \leqslant E + D$. If z' belongs to $\gamma \cap \widetilde{\gamma}$, then there is an edge $[z', v_V]$ in $\widehat{X}$ connecting z' to v_V, since we chose $A \geqslant D + E$ in Equation (5). If instead $z' \notin \gamma$, then z' lies along a de-electrified segment for some subspace $Z \in \mathcal{H}$ with $v_Z \in \gamma$; here $v_Z = v_W$ if $Z = \mathcal{N}_A(\rho^W_S)$ for some $W \in \mathfrak{S} - \{S\}$. Since Z is $\max\{K, 3E\}$–quasiconvex, there exists $z'' \in Z$ with $\mathsf{d}_X(z', z'') \leqslant \max\{K, 3E\}$. The situation is depicted in Figure 3, and again by our choice of A in Equation (5) there are edges $[v_Z, z'']$ and $[z'', v_V]$ in $\widehat{X}$. □

Lemma 4.23. *Let $x \in G$ and $U \in \mathfrak{S} - \{S\}$ be such that $\{x, U\}$ are minimal, and let $V \in \mathfrak{S} - \{S\}$ be such that $\mathsf{d}_X(\rho^U_S, \rho^V_S) \leqslant 2E$; we include the case $U = V$. Let $Y \in \mathcal{Y}$ be such that v_Y lies on a geodesic γ from x to v_U in $\widehat{X}$, and let $x' \in \gamma$ be the vertex of γ before v_Y. Then $\mathsf{d}_V(x, x') \leqslant E$.*

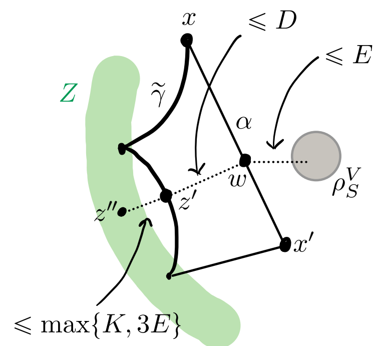

Figure 3: Proof of Lemma 4.22, in the case z' does not lie on γ.

Proof. Apply Lemma 4.22 to x, x', and the subsegment of γ between them. As in the proof, this produces some z, belonging to either $\gamma \cap \widetilde{\gamma}$ or some $Z \in \mathcal{H}$ where $v_Z \in \gamma$, such that $\mathsf{d}_X(z, \rho_S^V) \leqslant D + E + \max\{K, 3E\}$. By assumption, we have $\mathsf{d}_X(\rho_S^U, \rho_S^V) \leqslant 2E$, so $\mathsf{d}_X(z, \rho_S^U) \leqslant D + 4E + \max\{K, 3E\} \leqslant A$. As above, it follows that v_U is in the 2-neighborhood of the subsegment of γ between x and x'. However, the subpath of γ from x' to v_U has length at least 3, since it must contain v_Y and v_U, which are distance at least 2 apart. Hence we contradict the fact that γ is a geodesic between x and v_U. □

Proposition 4.24. *Let $x \in G$, $x' \in \bar{x}$, and $U \in \mathfrak{S} - \{S\}$ be such that $\{x, U\}$ and $\{x', U\}$ are minimal, and let $V \in \mathfrak{S} - \{S\}$ be such that $\mathsf{d}_X(\rho_S^U, \rho_S^V) \leqslant 2E$, where we allow $U = V$. Then*

$$\mathsf{d}_V(x, x') \leqslant 2\aleph + 9E,$$

where $\aleph$ is the constant from Remark 4.17.

Proof. We may assume $x \neq x'$, as otherwise there is nothing to prove. Let γ, γ' be $\widehat{X}$–geodesics from x and x' to v_U, respectively. Since $x, x' \in \overline{x}$, there is some $h \in N$ such that $hx = x' \neq x$.

Let $x_0 = x$ and $x_0' = x'$, and similarly let $\gamma_0 = \gamma$ and $\gamma_0' = \gamma'$. By Proposition 3.8 there exists a shortening pair (Y, h_Y) such that $h_Y h \prec h$ and $\mathsf{d}_Y^\pi(x_0, x_0') > L'/10$. By the triangle inequality, one of $\mathsf{d}_Y^\pi(x_0, v_U)$ and $\mathsf{d}_Y^\pi(v_U, x_0')$ is at least $L'/20$; we focus on the case $\mathsf{d}_Y^\pi(x_0, v_U) > L'/20$, as the other is dealt with analogously. Since $L'/20$ is greater than the constant C from Lemma 4.15, the cone point v_Y lies on γ_0. Now bend γ_0 at v_Y by h_Y^{-1}, and call this γ_1 with endpoint $x_1 = h_Y^{-1} x_0$. Then set $\gamma_1' = \gamma_0'$ and $x_1' = x_0'$. Notice that x_1 and x_1' differ by hh_Y, which is an N-conjugate of $h_Y h$ and so still has complexity strictly less than that of h. Moreover, both geodesics still have v_U as an endpoint.

Repeat the argument with x_1, x_1' and the associated geodesics. We proceed inductively until $x_k = x_k'$ for some $k \in \mathbb{N}$; we call this point x''. In other words, we eventually produce two geodesics γ_k and γ_k', both with endpoints v_U and $x'' \in \overline{x}$. Notice that x'' and v_U are minimal, as their distance is still the length of γ. Since γ_k is obtained by successively bending γ while fixing the endpoint v_U, there exists $v_Y \in \gamma \cap \gamma_k$. Moreover, if $z \in \gamma$ is the vertex before v_Y, then there exists $n \in N$ such that the vertex of γ_k before v_Y is nz. By Lemma 4.20, we can actually choose $n \in \mathrm{Stab}_G\, v_Y$, and the latter is H_Y by Corollary 3.6; hence $\mathsf{d}_V(z, nz) \leqslant \aleph$ by Remark 4.17. The situation is therefore as in Figure 4.

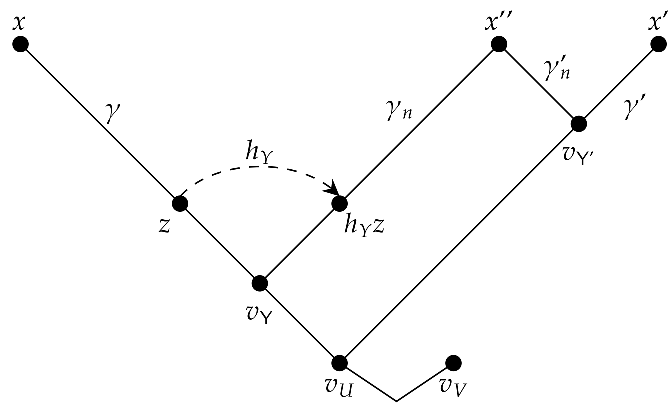

Figure 4: The geodesics from the proof of Proposition 4.24. After bending γ and γ' a finite number of times, their endpoints coincide. Notice that γ and γ_k must overlap on the segment of γ connecting v_U to the first vertex of the form v_Y for some $Y \in \mathcal{Y}$, as bending can occur only at cone points.

By Lemma 4.23 applied to the minimal pairs $\{x, U\}$ and $\{x'', U\}$, V, and Y, we obtain

$$\mathsf{d}_V(x, x'') \leqslant \mathsf{d}_V(x, z) + \mathsf{diam}\pi_V(z) + \mathsf{d}_V(z, nz) + \mathsf{diam}\pi_V(nz) + \mathsf{d}_V(nz, x'') \leqslant \aleph + 4E.$$

Symmetrically, we obtain $\mathsf{d}_V(x'', x') \leqslant \aleph + 4E$, completing the proof as $\mathsf{diam}\pi_V(x'') \leqslant E$. □

A similar statement holds for pairs of minimal domains:

Proposition 4.25. *Let $V \in \mathfrak{S} - \{S\}$ and $U, U' \in \overline{U} \in \mathfrak{S}/N - \{\overline{S}\}$ be such that $\{U, V\}$ and $\{U', V\}$ are minimal.*

- *If U and V are not transverse then $U = U'$.*
- *Otherwise $\mathsf{d}_V(\rho_V^U, \rho_V^{U'}) \leqslant 2\aleph + 25E$.*

Proof. Suppose first that $\mathsf{d}_{\widehat{X}}(v_U, v_V) = 2$. This occurs, in particular, if U and V are not transverse, since $\mathsf{d}_X(\rho_S^U, \rho_S^V) \leqslant 2E \leqslant 2A$ by Lemma 4.3. Then the union of any two geodesics $[v_U, v_V] \cup [v_V, v_{U'}]$ does not contain any cone point v_Y for $Y \in \mathcal{Y}$, as v_V is only $\widehat{X}$-adjacent to points in X, and so Lemma 4.21 gives that $U = U'$. This proves the first bullet point.

For the second bullet, the above argument lets us assume that $\mathsf{d}_{\widehat{X}}(v_U, v_V) \geqslant 3$, which in particular means that $\mathsf{d}_X(\rho_S^U, \rho_S^V) > 2A$. First, we prove that $\mathsf{diam}\pi_V(\mathcal{N}_A(\rho_S^U)) \leqslant 5E$. To see this, let $p \in \mathcal{N}_A(\rho_S^U)$, and fix $q \in \rho_S^U$, so that $\mathsf{d}(p, q) \leqslant A + E$. If a geodesic $[p, q]$ in X passed E-close to ρ_S^V, then $\mathsf{d}_X(\rho_S^U, \rho_S^V) \leqslant A + 2E$, and this is at most $2A$ by Equation (5), contradicting our assumption. Thus the bounded geodesic image axiom for $\mathfrak{S}$ yields that $\mathsf{d}_V(p, q) \leqslant E$. Hence $\pi_V(p) \subseteq \mathcal{N}_{2E}(\pi_V(q))$, from which the claim follows.

Now let γ and λ be geodesics connecting v_U (resp. $v_{U'}$) to v_V, and let p (resp. q) be the first point of γ after v_U (resp. the first point of λ after $v_{U'}$). Since $U, U' \in \overline{U}$, we may shorten the concatenation as in Proposition 4.24 to produce two geodesics, γ'' and λ'', each joining v_V to a single point $v_{U''} \in \overline{v_U}$, and such that γ'' is obtained by bending γ while λ'' is obtained by bending λ. If p'' is the first point of γ'' after $v_{U''}$, then $p'' \in \overline{p}$, and, as in the proof of Proposition 4.24, we obtain $\mathsf{d}_V(p, p'') \leqslant \aleph + 4E$.

Now let $r \in X$ be such that $\mathsf{d}_X(r, \rho_S^U) \leqslant E$ and $\mathsf{d}_V(r, \rho_V^U) \leqslant E$, which exists by the partial realization axiom Definition 4.1.(10). Since $r, p \in \mathcal{N}_A(\rho_S^U)$, we have that

$$\mathsf{d}_V(\rho_V^U, p'') \leqslant E + \mathsf{d}_V(r, p'') \leqslant E + \mathsf{diam}\pi_V(N_A(\rho_S^U)) + \mathsf{d}_V(p, p'') \leqslant \aleph + 10E.$$

Symmetrically, there exists $q'' \in \mathcal{N}_A(\rho_S^{U''})$ such that $\mathsf{d}_V(\rho_V^{U'}, q'') \leqslant \aleph + 10E$, thus concluding the proof as $\mathsf{diam}\pi_V(\mathcal{N}_A(\rho_S^{U''})) \leqslant 5E$. □

4.5 The quotient structure and projections

We are now ready to define the relative HHG structure on G/N. Recall that, by Remark 4.9, we are assuming that X is a Cayley graph for G. In particular, if we fix a word metric d_G on G with respect to a finite generating set, the projection map $\pi_S \colon G \to X$ is bijective at the level of vertices.

Construction 4.26. Fix a finite generating set $\mathcal{T}$ for G which induces a word metric d_G, let $\overline{\mathcal{T}}$ be its image in G/N, and let $\mathsf{d}_{G/N}$ be the word metric on G/N induced by $\overline{\mathcal{T}}$. The relative HHG structure on G/N has the following components.

- *Index set:* $\mathfrak{S}/N$. Recall that we denote the image of $U \in \mathfrak{S}$ by $\overline{U} \in \mathfrak{S}/N$.
- *Hyperbolic spaces:* The top-level space $\mathcal{C}\overline{S}$ is $\overline{X}$ as defined in Section 4.4. For each $\overline{U} \in \mathfrak{S}/N - \{\overline{S}\}$, we set
$$\mathcal{C}\overline{U} = \left(\bigcup_{U \in \overline{U}} \mathcal{C}U\right)/N.$$
Notice that, for every $U \in \overline{U}$, the projection map $\mathcal{C}U \to \mathcal{C}\overline{U}$ is an isometry. Indeed, by Corollary 3.9 N acts freely on X', and in particular no non-trivial $n \in N$ can fix v_U (hence U).
- *Relations:* Every domain in $\mathfrak{S}/N$ nests into $\overline{S}$. The relation between $\overline{U}, \overline{V} \in \mathfrak{S}/N - \{\overline{S}\}$ is the same as the relation between any minimal pair $\{U, V\}$ with $U \in \overline{U}$ and $V \in \overline{V}$. In particular, $\overline{U}$ is $\sqsubseteq$-minimal in $\mathfrak{S}/N$ if and only if every representative $U \in \overline{U}$ is $\sqsubseteq$-minimal in $\mathfrak{S}$.
- *Projection maps:* We identify G with X via the projection π_S, and we consider it as a subspace of $\widehat{X}$ via the inclusion $i \colon X \hookrightarrow \widehat{X}$. This way, G/N can be seen as a subgraph of $\overline{X}$. For every $\overline{g} \in G/N$ define $\pi_{\overline{S}} \colon G/N \to \overline{X}$ by $\pi_{\overline{S}}(\overline{g}) \coloneqq \overline{g}$; moreover, given $\overline{U} \in \mathfrak{S}/N - \{S\}$ set
$$\pi_{\overline{U}}(\overline{g}) \coloneqq \left(\bigsqcup_{\substack{g \in \overline{g},\, U \in \overline{U} \\ \{U, g\} \text{ minimal}}} \pi_U(g)\right)/N.$$
- *Relative projections:* Set $\rho_{\overline{S}}^{\overline{V}} = \overline{v_V}$ for some (equivalently, any) $V \in \overline{V}$. Furthermore, given $\overline{U}, \overline{V} \in \mathfrak{S}/N - \{\overline{S}\}$ with $\overline{U} \pitchfork \overline{V}$ or $\overline{V} \subsetneq \overline{U}$, define
$$\rho_{\overline{U}}^{\overline{V}} \coloneqq \left(\bigsqcup_{\substack{U \in \overline{U}, V \in \overline{V} \\ \{U, V\} \text{ minimal}}} \rho_U^V\right)/N.$$

Remark 4.27. Let $U \subsetneq S$ and $g \in G$ be minimal. An immediate consequence of how the projections are defined is that

$$\pi_U(g) \subseteq \pi_{\overline{U}}(\overline{g}).$$

Similarly, if $U, V \subsetneq S$ are minimal and non-orthogonal, then

$$\rho_U^V \subseteq \rho_{\overline{U}}^{\overline{V}}.$$

These convenient facts are useful for comparing the hierarchical structure for $(G, \mathfrak{S})$ and that of $(G/N, \mathfrak{S}/N)$.

Remark 4.28. Let $\overline{U}, \overline{V} \in \mathfrak{S}/N - \{\overline{S}\}$, and fix a representative $U \in \overline{U}$. If there exists $V \in \overline{V}$ such that $\mathsf{d}_X(\rho_S^U, \rho_S^V) \leqslant 2A$ (for example if U and V are not transverse, by Lemma 4.3), then there can be at most one such domain $V \in \overline{V}$, as we argued in the proof of Proposition 4.25. As a consequence, for every other $V' \in \overline{V}$ we have that $\mathsf{d}_X(\rho_S^U, \rho_S^{V'}) > 2A > E$, so $V' \pitchfork U$. This proves that the relations $\sqsubseteq$ and $\perp$ in $\mathfrak{S}/N$ are well-defined, as either there is a unique minimal pair including U, or every pair $\{U, V\}$ is transverse. This also shows that $\overline{U}$ and $\overline{V}$ are orthogonal (resp. nested) if and only if they admit orthogonal (resp. nested) representatives U and V, as such representatives are minimal.

4.5.1 Lifting minimal configurations

Before we prove that the above construction yields a hierarchy structure on G/N, we record a few technical lemmas about projections and minimal collections. We first argue that certain configurations admit pairwise minimal representatives.

Lemma 4.29 (Lifting triples)**.** *Given $\overline{x}, \overline{y}, \overline{z} \in G/N$, there exist representatives $x \in \overline{x}, y \in \overline{y}$, and $z \in \overline{z}$ such that $\{x, y, z\}$ are pairwise minimal. Furthermore, the same holds when any of the elements of G/N are replaced with domains in $\mathfrak{S}/N - \{\overline{S}\}$.*

Proof. Fix $x \in \overline{x}$. Considering $\overline{x}, \overline{y}, \overline{z}$ as vertices of $\overline{X}$, pick geodesics in $\overline{X}$ to form a geodesic triangle. By Proposition 3.12 we can lift this to a geodesic triangle in $\widehat{X}$ based at $x \in \overline{x}$. The vertices of this triangle are pairwise minimal by construction. The same argument holds if we replace any vertex of the triangle in $\overline{X}$ with the image of the cone point over a domain in $\mathfrak{S} - \{S\}$, thus proving the "furthermore" part of the statement. □

Lemma 4.30. *Let $\overline{x}, \overline{y} \in G/N$ and $\overline{U}_1, \dots, \overline{U}_k \in \mathfrak{S}/N - \{\overline{S}\}$. There exist representatives $\{x, y, U_1, \dots, U_k\}$ such that $\{x, y, U_i\}$ are pairwise minimal for every $i \leqslant k$.*

Proof. We proceed by induction on k. The base case $k = 1$ is Lemma 4.29. Suppose that $\{x, y, U_1, \dots, U_{k-1}\}$ are as in the statement. There are $\{x', y', U_k\}$ pairwise minimal representatives by Remark 4.28, where $x' \in \overline{x}$ and $y' \in \overline{y}$. Up to the action of N, we can assume that $x = x'$. Let $h \in N$ map y to y'. For every $i \leqslant k-1$, consider a geodesic triangle with vertices $\{x, y, v_{U_i}\}$, and also consider a geodesic triangle with vertices $\{x, y', v_{U_k}\}$, as in Figure 5.

If $hy = y$, we are done. Otherwise, by Proposition 3.8 there exists a shortening pair (Y, h_Y), so that $\mathsf{d}_Y^\pi(y, y') > L'/10$. By triangle inequality, one of $\mathsf{d}_Y^\pi(y, x)$ and $\mathsf{d}_Y^\pi(x, y')$ is at least $L'/20$.

If $\mathsf{d}_Y^\pi(y, x) > L'/20$, then for every $i \leqslant k-1$, one of $\mathsf{d}_Y^\pi(y, v_{U_i})$ and $\mathsf{d}_Y^\pi(v_{U_i}, x)$ is at least $L'/40 > C$. Lemma 4.15 then implies that v_Y lies both on every geodesic $[x, v_{U_i}]$ and every geodesic $[x, y]$. In particular, v_Y is a cut vertex of the triangle with vertices $\{x, y, v_{U_i}\}$. For

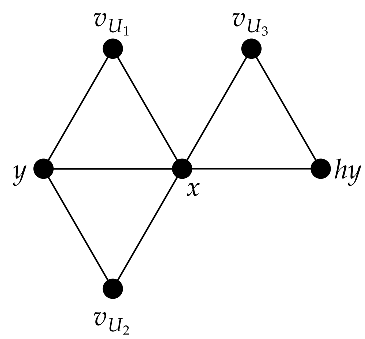

Figure 5: The various lifts from the proof of Lemma 4.30 (for $k = 3$).

the same reason, if $\mathsf{d}_Y^\pi(x, y') > L'/20$, then v_Y is a cut vertex of the triangle with vertices $\{x, y', v_{U_k}\}$. In either case, applying h_Y to the every vertex in the configuration from Figure 5 between v_Y and y' reduces the complexity of h, and we conclude by induction. □

If the given domains are all non-transverse to some domain, the lifts from Lemma 4.30 can be chosen to be pairwise minimal:

Lemma 4.31. *Let $\{\overline{U}_0, \ldots, \overline{U}_k\}$ be a collection of domains in $\mathfrak{S}/N - \{\overline{S}\}$ such that $\overline{U}_i$ and $\overline{U}_0$ are not transverse for any $i \neq 0$, and let $\overline{x}, \overline{y} \in G/N$. Then there exist pairwise minimal representatives $\{x, y, U_0, \ldots, U_k\}$.*

Proof. We first prove the existence of pairwise minimal representatives of the domains alone. Fix $U_0 \in \overline{U}_0$. Since $\overline{U}_i$ is not transverse to $\overline{U}_0$ for any i, by Remark 4.28 there is a unique representative $U_i \in \overline{U}_i$ such that $\{U_1, U_0\}$ are minimal, and $\rho_S^{U_i}$ at distance at most $2E$ from $\rho_S^{U_0}$ by Lemma 4.3. Since $\mathsf{d}_X(\rho_S^{U_i}, \rho_S^{U_j}) \leqslant 5E \leqslant 2A$, it follows from Remark 4.28 that $\{U_i, U_j\}$ are minimal for all $i \neq j$.

Next, we prove by induction on k that there exist minimal representatives $\{x, U_0, \ldots, U_k\}$. If $k = 0$, there is nothing to prove. Assume now that there exist pairwise minimal representatives $\{x, U_0', \ldots, U_{k-1}'\}$, and let $\{U_0, \ldots, U_k\}$ be pairwise minimal representatives, whose existence is guaranteed by the above argument. Up to the action of N we can assume that $U_0 = U_0'$, and by Proposition 4.25 this implies that $U_i = U_i'$ for every $i \leqslant k-1$ (but possibly x and U_k are not yet minimal). Let $x' \in \overline{x}$ be minimal with respect to U_k. Consider geodesics from x to every v_{U_i} for $i \leqslant k-1$, from v_{U_i} to v_{U_j} for every i, j, and from v_{U_k} to x', as in Figure 6. Notice that each such geodesic lifts a geodesic of $\overline{X}$, as its endpoints are minimal.

Let $h \in N$ be such that $hx = x'$. We proceed by induction on the complexity of h. If $hx = x$ we are done, as then x and U_k are already minimal. Otherwise, there is (Y, h_Y) a shortening pair by Proposition 3.8. If $\mathsf{d}_Y^\pi(v_{U_k}, x') > C$ then v_Y lies on the geodesic between v_{U_k} and x'. In this case, we bend this geodesic at v_Y, and conclude by induction since $h_Y x'$ differs from x by $h_Y h \prec h$. If instead $\mathsf{d}_Y^\pi(v_{U_k}, x') \leqslant C$, then for every $i \leqslant k-1$ we have that

$$\mathsf{d}_Y^\pi(x, v_{U_i}) \geqslant \mathsf{d}_Y^\pi(x, x') - \mathsf{d}_Y^\pi(v_{U_i}, v_{U_k}) - \mathsf{d}_Y^\pi(v_{U_k}, x') \geqslant L'/10 - 2C > C.$$

Here, we used that, since U_i and U_k are not transverse, they lie at distance 2 in $\widehat{X}$, so any $\widehat{X}$–geodesic connecting them belongs to X' and cannot pass through v_Y. Hence v_Y lies on

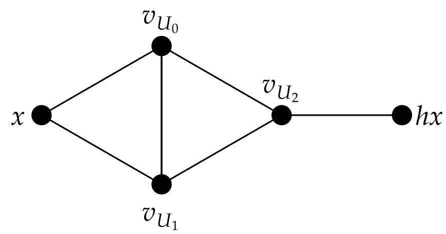

Figure 6: The various lifts from the first part of the proof of Lemma 4.31 (with $k = 2$). Each edge in the picture is a geodesic with whose endpoints are minimal.

the geodesic between x and each v_{U_i}, so we bend each of these geodesics at v_Y by applying h_Y to the geodesic connecting v_Y to x'. Again $h_Y x'$ differs from x by $h_Y h \prec h$, and we conclude by induction.

Finally, let $\{x, U_1, \ldots, U_k\}$ and $\{y, U'_1, \ldots, U'_k\}$ be pairwise minimal collections, which exist by the above argument. As before, up to N-translation we can actually assume that $U_i = U'_i$ for every i. Consider geodesics from x to v_{U_i} for every i, from v_{U_i} to v_{U_j} for every $i \neq j$, from v_{U_i} to y for every i, and from y to some $x' \in \overline{X}$ such that $\{y, x'\}$ are minimal. The situation is as in Figure 7.

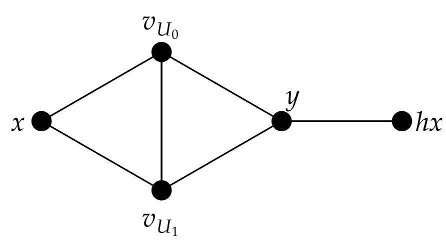

Figure 7: The various lifts from the final part of the proof of Lemma 4.31 (with $k = 1$).

The argument to find pairwise minimal representatives is now very similar to the one with x alone. Let $h \in N$ map x to x'. If $hx = x$ we are done; otherwise let (Y, h_Y) be a shortening pair. If $\mathsf{d}^{\pi}_Y(y, x') > C$ then v_Y lies on the geodesic between y and x'; in this case we bend this geodesic at v_Y and conclude by induction. If instead $\mathsf{d}^{\pi}_Y(y, x') \leqslant C$ then $\mathsf{d}^{\pi}_Y(x, y) \geqslant L'/10 - C > 3C$, where we used Equation (7). In turn, for every $i, j \leqslant k$ one of $\mathsf{d}^{\pi}_Y(x, v_{U_i})$ and $\mathsf{d}^{\pi}_Y(v_{U_j}, y)$ is greater than C, as otherwise we would have that

$$\mathsf{d}^{\pi}_Y(x, y) \leqslant \mathsf{d}^{\pi}_Y(x, v_{U_i}) + \mathsf{d}^{\pi}_Y(v_{U_i}, v_{U_j}) + \mathsf{d}^{\pi}_Y(v_{U_j}, y) \leqslant 3C.$$

In particular, this means that either $\mathsf{d}^{\pi}_Y(x, v_{U_i}) > C$ for every i, or the same holds with y replacing x. In other words, by Lemma 4.15, v_Y lies on either every geodesic coming out of x, or every geodesic ending at y. In either case v_Y is a cut point of the configuration from Figure 7, so we can apply h_Y to every point between v_Y and x' and again conclude by induction. □

4.5.2 Bounded projections

Our next goal is to provide some bounds on the projections from Construction 4.26. We first show that the new maps π_* and ρ_*^* send points to uniformly bounded sets.

Proposition 4.32. *Let $\beth = 2\aleph + 27E$, where $\aleph$ is the constant from Remark 4.17. For any $\overline{x} \in G/N$ and $\overline{U} \in \mathfrak{S}/N$, the projection $\pi_{\overline{U}}(\overline{x})$ has diameter at most $\beth$ in $\mathcal{C}\overline{U}$. Analogously, if $\overline{U}, \overline{V} \in \mathfrak{S}/N$ satisfy $\overline{U} \pitchfork \overline{V}$ or $\overline{V} \sqsubseteq \overline{U}$, then $\rho_{\overline{U}}^{\overline{V}}$ has diameter bounded by $\beth$ in $\mathcal{C}\overline{U}$.*

Proof. If $\overline{U} = \overline{S}$, then $\pi_{\overline{S}}(\overline{x}) = \overline{x}$, seen as a point in $\overline{X}$, and we have nothing to prove. Thus suppose $\overline{U} \neq \overline{S}$. Since being minimal is an N–equivariant relation, it suffices to consider a fixed $U \in \overline{U}$, so that we can identify $\mathcal{C}\overline{U}$ with $\mathcal{C}U$. By Proposition 4.24, if $x, x' \in \overline{x}$ are both minimal with U, then $\mathsf{d}_U(x, x') \leqslant 2\aleph + 9E$, so $\mathsf{diam}(\pi_U(x) \cup \pi_U(x')) \leqslant 2\aleph + 11E \leqslant \beth$.

For the second statement, if $\overline{U} = \overline{S}$ then $\rho_{\overline{S}}^{\overline{V}} = \overline{v_V}$ is a point. Otherwise fix $U \in \overline{U}$, and suppose $\{V', U\}$ and $\{V, U\}$ are both minimal, with $V, V' \in \overline{V}$. The result follows by applying Proposition 4.25 and using that ρ_U^V, $\rho_U^{V'}$ have diameter at most E. ☐

Lemma 4.33. *Let $x, y \in X$, and suppose $U \in \mathfrak{S} - \{S\}$. If $\{U, x, y\}$ are pairwise minimal, then*

$$\mathsf{d}_U(x, y) - 2\beth \leqslant \mathsf{d}_{\overline{U}}(\overline{x}, \overline{y}) \leqslant \mathsf{d}_U(x, y) \quad \textit{and} \quad \mathsf{d}_{\overline{X}}(\overline{x}, \overline{y}) \leqslant \mathsf{d}_X(x, y).$$

Moreover, if $\overline{V} \in \mathfrak{S}/N - \{\overline{S}\}$ satisfies $\overline{U} \pitchfork \overline{V}$ or $\overline{V} \subsetneq \overline{U}$ and $\{x, U, V\}$ are pairwise minimal representatives, then

$$\mathsf{d}_U(x, \rho_U^V) - 2\beth \leqslant \mathsf{d}_{\overline{U}}(\overline{x}, \rho_{\overline{U}}^{\overline{V}}) \leqslant \mathsf{d}_U(x, \rho_U^V) \quad \textit{and} \quad \mathsf{d}_{\overline{X}}(\overline{x}, \overline{v_V}) \leqslant \mathsf{d}_X(x, \rho_S^V) + 1.$$

Proof. For any pair $\{x, y\}$, we have that $\mathsf{d}_{\overline{X}}(\overline{x}, \overline{y}) \leqslant \mathsf{d}_{\widehat{X}}(x, y) \leqslant \mathsf{d}_X(x, y)$. When $y \in \overline{y}$ is replaced with $V \in \overline{V} \neq \overline{S}$, we have $\mathsf{d}_{\overline{X}}(\overline{x}, \overline{v_V}) \leqslant \mathsf{d}_{\widehat{X}}(x, v_V) \leqslant \mathsf{d}_X(x, \rho_S^V) + 1$.

Furthermore, by definition of the projections $\pi_{\overline{U}}$ and the fact that $\mathcal{C}\overline{U}$ is identified with $\mathcal{C}U$, we have $\pi_{\overline{U}}(\overline{x}) \supseteq \pi_U(x)$ by Remark 4.27, so $\mathsf{d}_U(x, y) \geqslant \mathsf{d}_{\overline{U}}(\overline{x}, \overline{y})$. On the other hand, by Proposition 4.32, $\pi_{\overline{U}}(\overline{x})$ and $\pi_{\overline{U}}(\overline{y})$ have diameter bounded by $\beth$, so $\mathsf{d}_{\overline{U}}(\overline{x}, \overline{y}) \geqslant \mathsf{d}_U(x, y) - 2\beth$. Analogously, $\rho_{\overline{U}}^{\overline{V}}$ contains ρ_U^V and has diameter at most $\beth$, so the second statement follows analogously. ☐

Finally, we prove that if $x \in G$ and $U \in \mathfrak{S}$ are "almost" minimal, then the projection of x to U is almost the projection of a minimal element:

Lemma 4.34. *Let $x \in X$, $y \in \widehat{X}$ and $U \in \mathfrak{S} - \{S\}$. Suppose that $\{x, y\}$ are minimal and $\mathsf{d}_{\widehat{X}}(v_U, y) \leqslant 2$. Then there exists $x^* \in \overline{x}$ such that $\{x^*, U\}$ are minimal and*

$$\mathsf{d}_U(x, x^*) \leqslant 9\aleph + 28E.$$

Proof. If $\{x, U\}$ are already minimal, there is nothing to prove, so we henceforth assume the contrary. Let $\gamma = [y, x]$ be an $\widehat{X}$–geodesic, and up to replacing y by some point on γ, we can assume that $\gamma \cap \mathcal{N}_2(v_U) = \{y\}$ (notice that $\mathsf{d}_{\widehat{X}}(x, v_U) \geqslant 2$, or they would be minimal). If $\eta = [v_U, y]$ is an $\widehat{X}$–geodesic, then y is the only point of η which can be of the form v_Y for some $Y \in \mathcal{Y}$, because the link of v_U in $\widehat{X}$ belongs to X.

We now describe an algorithm to find the required $x^* \in \overline{x}$. To initialize the procedure, set $\gamma_0 = \gamma$ and $x_0 = x$; furthermore, let $\widetilde{x}_0 \in \overline{x}$ be minimal with U, let σ_0 be an $\widehat{X}$-geodesic connecting v_U to $\widetilde{x}_0$, and let $h_0 \in N$ map x_0 to $\widetilde{x}_0$.

Suppose we are given geodesics $\gamma_i = [y, x_i]$ and $\sigma_i = [v_U, \widetilde{x}_i]$, where $x_i, \widetilde{x}_i \in \overline{x}$ and the endpoints of both geodesics are minimal. Suppose $h_i \in N$ maps x_i to $\widetilde{x}_i$. If $\widetilde{x}_i = x_i$, then set $x^* = x_i$ and stop the algorithm. Otherwise, let (Y, h_Y) be a shortening pair, as in Proposition 3.8. Then v_Y must lie on one of σ_i or γ_i since $L'/10 > 3C$; notice that it cannot lie in the interior of η, which is a single point in X.

1. If $v_Y \in \sigma_i$, let $\widetilde{x}_{i+1} = h_Y \widetilde{x}_i$, and let $\sigma_{i+1} = [v_U, \widetilde{x}_{i+1}]$ be obtained by bending σ_i at v_Y by h_Y. Set $\gamma_{i+1} = \gamma_i$, $x_{i+1} = x_i$, and $h_{i+1} = h_Y h_i \prec h_i$. We now repeat the procedure with the data indexed by $i+1$.
2. Suppose instead that $v_Y \in \gamma_i$. Let $x_{i+1} = h_Y^{-1} x_i$, and let $\gamma_{i+1} = [y, x_{i+1}]$ be obtained by bending γ_i at v_Y by h_Y^{-1}. There are two sub-cases to consider.
   (a) If $\gamma_{i+1} \cap \mathcal{N}_2(v_U) = \{y\}$, we set $\eta_{i+1} = \eta_i$, $\widetilde{x}_{i+1} = \widetilde{x}_i$, and $h_{i+1} = h_i h_Y$, which is conjugate to $h_Y h_i \prec h_i$ and therefore has the same complexity. We now repeat the procedure with the data indexed by $i+1$.
   (b) If instead $\gamma_{i+1} \cap \mathcal{N}_2(v_U)$ contains some other point z, we stop the algorithm.

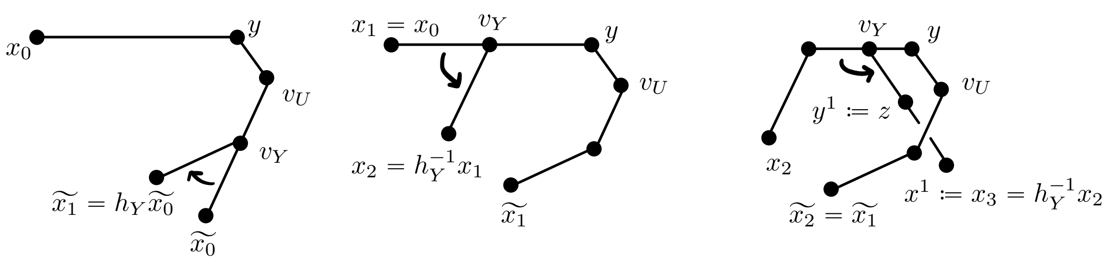

Figure 8: An example of the algorithm in the proof of Lemma 4.34. From left to right, we applied step 1, then step 2(a), and finally step 2(b), at which point we reached a bad ending and stopped the algorithm. We would next apply the algorithm to the points x^1 and y^1.

Good ending. When running the above algorithm, if we never encounter the termination condition from Item 2b, then each step reduces the complexity of h_i. Since the complexity is a good ordering, we must eventually find some $n \in \mathbb{N}$ such that $x^* = x_n = \widetilde{x}_n$, which is minimal with U. Since γ_n is obtained by successively bending γ_0, there exists $Y \in \mathcal{Y}$ such that $v_Y \in \gamma_0 \cap \gamma_n$. Let $t \in \gamma_0$ be the vertex after v_Y, so that the first vertex of γ_n after v_Y is of the form $h_Y t$ for some $h_Y \in H_Y$ (to see this, we can argue exactly as in the proof of Proposition 4.24). Then

$$\mathsf{d}_U(x, x^*) \leqslant \mathsf{d}_U(x, t) + \mathsf{diam}\pi_U(t) + \mathsf{d}_U(t, h_Y t) + \mathsf{diam}\pi_U(h_Y t) + \mathsf{d}_U(h_Y t, x^*).$$

The third term is at most $\aleph$ by Remark 4.17. Moreover the segment of γ_0 between t and x does not pass 2-close to v_U, so Lemma 4.22 gives that the first term is bounded by E. Similarly, since we never encountered the termination condition from Item 2b, the last term is also bounded by E. Therefore

$$\mathsf{d}_U(x, x^*) \leqslant \aleph + 4E, \tag{9}$$

which satisfies the requirement of the statement.

Bad ending. Suppose instead that the algorithm terminated because some $z \in \gamma_{i+1} - \{y\}$ is within distance 2 from v_U. We can assume that $\mathsf{d}_{\widehat{X}}(z, v_U) = 2$. Notice that the above argument applies to x and x_i, giving that $\mathsf{d}_U(x, x_i) \leqslant \aleph + 4E$. Moreover, since x_{i+1} differs from x_i by an element in some H_Y, we have that

$$\mathsf{d}_U(x, x_{i+1}) \leqslant \mathsf{d}_U(x, x_i) + \mathsf{diam}\pi_U(x_i) + \mathsf{d}_U(x_i, x_{i+1}) \leqslant 2\aleph + 5E. \tag{10}$$

Now set $x^1 = x_{i+1}$ and $y^1 = z$, and notice that $\mathsf{d}_{\widehat{X}}(y^1, x^1) < \mathsf{d}_{\widehat{X}}(y, x)$. If we repeat the whole procedure with $\{x^1, y^1, U\}$, we either obtain a good ending or find some $\{x^2, y^2\}$ with $\mathsf{d}_{\widehat{X}}(y^2, x^2) < \mathsf{d}_{\widehat{X}}(y^1, x^1)$, and so on. We cannot keep falling into the termination condition from Item 2b more than four times in total. Indeed, if so, then we would have

$$\mathsf{d}_{\widehat{X}}(y^5, x^5) < \mathsf{d}_{\widehat{X}}(y^4, x^4) < \ldots < \mathsf{d}_{\widehat{X}}(x, y),$$

and so $\mathsf{d}_{\widehat{X}}(y^5, x^5) \leqslant \mathsf{d}_{\widehat{X}}(x, y) - 5$. In turn,

$$\mathsf{d}_{\overline{X}}(\overline{v_U}, \overline{x}) \leqslant \mathsf{d}_{\widehat{X}}(v_U, x^5) \leqslant 2 + \mathsf{d}_{\widehat{X}}(y^5, x^5) \leqslant \mathsf{d}_{\widehat{X}}(x, y) - 3.$$

However this would be a contradiction, since

$$\mathsf{d}_{\overline{X}}(\overline{v_U}, \overline{x}) \geqslant \mathsf{d}_{\overline{X}}(\overline{x}, \overline{y}) - \mathsf{d}_{\overline{X}}(\overline{v_U}, \overline{y}) \geqslant \mathsf{d}_{\widehat{X}}(x, y) - 2,$$

where we used that $\{x, y\}$ are minimal and that $\mathsf{d}_{\overline{X}}(\overline{v_U}, \overline{y}) \leqslant \mathsf{d}_{\widehat{X}}(v_U, y) \leqslant 2$.

Since the process must have a good ending after at most four bad endings, there exists $j \leqslant 4$ and some $x^* \in \overline{x}$ which is minimal with U, such that $\mathsf{d}_U(x^j, x^*) \leqslant \aleph + 4E$. We then use (9) and (10) to conclude that

$$\mathsf{d}_U(x, x^*) \leqslant \mathsf{d}_U(x, x^1) + \mathsf{diam}\pi_U(x^1) + \ldots + \mathsf{d}_U(x^{j-1}, x^j) + \mathsf{diam}\pi_U(x^j) + \mathsf{d}_U(x^j, x^*) \leqslant$$

$$\leqslant 4(2\aleph + 5E + E) + \aleph + 4E = 9\aleph + 28E,$$

as required. □

4.6 The quotient satisfies the HHG axioms

We are now ready to show that G/N is a relative HHG. More precisely, we shall check that each axiom from Definition 4.1 is satisfied by the structure described in Construction 4.26 for some choice of the hierarchy constant. We shall then set $\overline{E}$ to be the maximum of these constants, thus ensuring that all axioms hold. We advise the reader to keep the list of all constants from Appendix A at hand throughout the proof.

Proof of Theorem 4.13. We check each axiom in turn.

(1) Projections: The maps $\pi_{\overline{U}}$ for $\overline{U} \in \mathfrak{S}/N$ send points to uniformly bounded diameter sets by Proposition 4.32. We now check that $\pi_{\overline{U}}$ is coarsely Lipschitz with respect to the fixed word metric $\mathsf{d}_{G/N}$ induced by d_G. Given $\overline{g}, \overline{h} \in G/N$ such that $\mathsf{d}_{G/N}(\overline{g}, \overline{h}) = 1$, we will uniformly bound $\mathsf{d}_{\overline{U}}(\overline{g}, \overline{h})$. Let $g, h \in G$ be such that $\mathsf{d}_G(g, h) = \mathsf{d}_{G/N}(\overline{g}, \overline{h}) = 1$, so that for every $V \in \mathfrak{S}$ we have that $\mathsf{d}_V(g, h) \leqslant 2E$, as projections in $(G, \mathfrak{S})$ are E-coarsely-Lipschitz. If $\overline{U} = \overline{S}$ then

$$\mathsf{d}_{\overline{X}}(\pi_{\overline{S}}(\overline{g}), \pi_{\overline{S}}(\overline{g})) = \mathsf{d}_{\overline{X}}(\overline{g}, \overline{h}) \leqslant \mathsf{d}_X(g, h) \leqslant 2E.$$

Now suppose $\overline{U} \neq \overline{S}$, and let $\overline{T} \subseteq \overline{X}$ be a geodesic triangle with vertices $\{\overline{g}, \overline{h}, \overline{v_U}\}$. Lift $\overline{T}$ to a geodesic triangle T in $\widehat{X}$ with vertices $\{g, h', v_U\}$, where $h' \in \overline{h}$, and consider

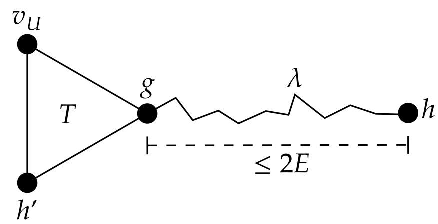

Figure 9: The configuration in the proof of the Projection axiom. In the picture, straight lines are $\widehat{X}$-geodesics, while the zig-zag path λ is an X–geodesic of length at most $2E$.

an X-geodesic λ from g to h, whose length (in X, hence also in $\widehat{X}$) is at most $2E$. The situation is as in Figure 9.

Notice that g, h', v_U are minimal, so $\mathsf{d}_{\overline{U}}(\overline{g}, \overline{h}) \leqslant \mathsf{d}_U(\pi_U(g), \pi_U(h'))$. If $h = h'$, then $\mathsf{d}_U(\pi_U(g), \pi_U(h')) = \mathsf{d}_U(\pi_U(g), \pi_U(h)) \leqslant 2E$, and we are done. Otherwise we proceed by induction on the complexity of the element $n \in N$ mapping h to h'. By Proposition 3.8, there exists a shortening pair (Y, h_Y). If $\mathsf{d}^\pi_Y(h', g) \leqslant 2C$ then $\mathsf{d}^\pi_Y(h', h) \leqslant 2C + 2EJ$, as every two consecutive points of λ are X–adjacent and π_Y is J Lipschitz. But then Equation (7) gives that $\mathsf{d}^\pi_Y(h', h) \leqslant L/10$, contradicting the definition of a shortening pair. Thus we must have that $\mathsf{d}^\pi_Y(h', g) > 2C$, so one of $\mathsf{d}^\pi_Y(h', v_U)$ and $\mathsf{d}^\pi_Y(v_U, g)$ is at least C by the triangle inequality. In other words, v_Y is a cut vertex for T by Lemma 4.15. Apply h_Y to the connected component of $T - \{v_Y\}$ containing h'. After this procedure, g, the image of v_U, and $h_Y h'$ are again minimal, and now $h_Y h'$ differs from h by $h_Y n \prec n$. Proceeding inductively, we eventually find some $U' \in \overline{U}$ such that $\{g, h, v_{U'}\}$ are minimal and $\mathsf{d}_{U'}(g, h) = \mathsf{d}_U(g, h') \leqslant 2E$, so $\mathsf{d}_{\overline{U}}(\overline{g}, \overline{h}) \leqslant \mathsf{d}_{U'}(g, h) \leqslant 2E$.

(2) Nesting: The nesting relation $\sqsubseteq$ and the subsets $\rho^{\overline{V}}_{\overline{W}}$ for $\overline{V} \subsetneq \overline{W}$ are defined in Construction 4.26. These sets have uniformly bounded diameter by Proposition 4.32.

(3) Finite complexity: If $\overline{U}_1 \sqsubseteq \overline{U}_2 \sqsubseteq \ldots$ then by Lemma 4.31 we may choose pairwise minimal representatives $U_1 \sqsubseteq U_2 \sqsubseteq \ldots$ of the $\overline{U}_i$. Finite complexity in $(G, \mathcal{S})$ then implies that both sequences are eventually constant.

(4) Orthogonality: The relation $\perp$ is defined in Construction 4.26, and by construction has the same properties as $\perp$ in $(G, \mathfrak{S})$. Thus the axiom holds by the orthogonality axiom in $(G, \mathfrak{S})$.

(5) Containers: To show containers exist, let $\overline{T} \in \mathfrak{S}/N$, and consider $\overline{U} \in (\mathfrak{S}/N)_{\overline{T}}$ such that $\overline{\mathcal{V}} := \{\overline{V} \in (\mathfrak{S}/N)_{\overline{T}} \mid \overline{V} \perp \overline{U}\} \neq \varnothing$.

If $\overline{T} = \overline{S}$, let $U \in \overline{U}$, and let W be the container for U in S. Now every $\overline{V}$ which is orthogonal to $\overline{U}$ has a representative V which is orthogonal to U and therefore nested in W. Then $\overline{V} \sqsubseteq \overline{W}$ by Remark 4.28, proving that $\overline{W}$ is the container for $\overline{U}$ in $\overline{S}$.

Now assume that $\overline{T} \neq \overline{S}$, and choose $\{U, T\}$ minimal representatives of $\overline{U}$ and $\overline{T}$. Since $U \sqsubseteq T$, we have $\mathsf{d}_X(\rho^U_S, \rho^T_S) \leqslant E$ by consistency in $\mathfrak{S}$. By Remark 4.28, each $\overline{V} \in \overline{\mathcal{V}}$ has a unique representative V which is orthogonal to U. Let $\mathcal{V}$ be the collection of such representatives. Notice that $\mathsf{d}_X(\rho^V_S, \rho^U_S) \leqslant 2E$ by Lemma 4.3. Combining this with the previous inequality yields that $\mathsf{d}_X(\rho^V_S, \rho^T_S) \leqslant 4E \leqslant 2A$, so V and T are minimal. Since

$\overline{V} \sqsubseteq \overline{T}$, the minimal representatives V and T must be nested. Now let W be the container for U in T, which contains every $V \in \mathcal{V}$. It then follows from Remark 4.28 that $\overline{W} \sqsubsetneq \overline{T}$ and $\overline{W}$ contains every $\overline{V} \in \overline{\mathcal{V}}$, as required.

(6) Transversality: The relation $\pitchfork$ and the sets $\rho^{\overline{U}}_{\overline{W}} \subseteq \mathcal{C}\overline{W}$ and $\rho^{\overline{W}}_{\overline{U}} \subseteq \mathcal{C}\overline{U}$ are defined in Construction 4.26. These sets have uniformly bounded diameter by Proposition 4.32.

(7) Consistency: Let $\overline{x} \in G/N$, and let $\overline{U}, \overline{V} \in \mathfrak{S}/N$ satisfy $\overline{U} \pitchfork \overline{V}$. By Lemma 4.29, we may choose $x \in \overline{x}$, $U \in \overline{U}$, and $V \in \overline{V}$ such that $\{x, U, V\}$ are pairwise minimal. By the definition of the relation $\pitchfork$ in $\mathfrak{S}/N$, we have $U \pitchfork V$. It follows from the consistency axiom in $(G, \mathfrak{S})$ that

$$\min\{\mathsf{d}_U(x, \rho^V_U), \mathsf{d}_V(x, \rho^U_V)\} \leqslant E.$$

The first consistency inequality now follows immediately from Lemma 4.33.

For the second statement, suppose that $\overline{U} \sqsubseteq \overline{V}$ and either $\overline{V} \sqsubsetneq \overline{W}$, or $\overline{V} \pitchfork \overline{W}$ and $\overline{W} \not\perp \overline{U}$. Suppose first that $\overline{W} \neq \overline{S}$. We may choose $U \in \overline{U}$, $V \in \overline{V}$, and $W \in \overline{W}$ with $\{U, V, W\}$ pairwise minimal by Lemma 4.29, and the same relations hold between U, V, and W as between $\overline{U}, \overline{V}, \overline{W}$. It follows from the consistency axiom in $(G, \mathfrak{S})$ that ρ^U_W is coarsely equal to ρ^V_W in $\mathcal{C}W = \mathcal{C}\overline{W}$, and the result again follows from Lemma 4.33.

If instead $\overline{W} = \overline{S}$ let $U \in \overline{U}$ and $V \in \overline{V}$ be nested representatives. Then $\mathsf{d}_X(\rho^U_S, \rho^V_S) \leqslant E$ by consistency in $\mathfrak{S}$. As $E \leqslant 2A$, we have $\mathsf{d}_{\overline{X}}(\overline{v_U}, \overline{v_V}) \leqslant \mathsf{d}_{\widehat{X}}(v_U, v_V) \leqslant 2$.

(8) Hyperbolicity: If $(G, \mathfrak{S})$ is a HHG, then $\mathcal{C}\overline{U}$ is hyperbolic for all $U \in \mathfrak{S}$ by definition if $U \neq S$ and by Theorem 3.13 if $U = S$. If $(G, \mathfrak{S})$ is only a relative HHG, then the spaces associated to all domains that are not $\sqsubseteq$–minimal in $\mathfrak{S}/N$ are E–hyperbolic for the same reason.

(9) Bounded geodesic image: Let $\overline{W} \in \mathfrak{S}/N$, let $\overline{V} \sqsubsetneq \overline{W}$, and let $\overline{x}, \overline{y} \in G/N$ satisfy $\mathsf{d}_{\overline{V}}(\overline{x}, \overline{y}) > E$. Let $\overline{\gamma}$ be a geodesic in $\mathcal{C}\overline{W}$ from $\pi_{\overline{W}}(\overline{x})$ to $\pi_{\overline{W}}(\overline{y})$. We will show that a uniform neighborhood of $\rho^{\overline{V}}_{\overline{W}}$ intersects $\overline{\gamma}$.

Suppose first that $\overline{W} \neq \overline{S}$, so that $\mathcal{C}\overline{W} = \mathcal{C}W$. Since $\overline{V} \sqsubsetneq \overline{W}$, Lemma 4.31 implies that there exist pairwise minimal representatives $\{x, y, V, W\}$. Since $\{x, V\}$ are minimal and $\mathcal{C}\overline{V} = \mathcal{C}V$, we have $\pi_V(x) \subseteq \pi_{\overline{V}}(\overline{x})$, and similarly for y. Thus $\mathsf{d}_V(x, y) \geqslant \mathsf{d}_{\overline{V}}(\overline{x}, \overline{y}) > E$. By Remark 4.27 we have $\rho^V_W \subseteq \rho^{\overline{V}}_{\overline{W}}$, so the bounded geodesic image axiom in $(G, \mathfrak{S})$ provides a point $w \in \mathcal{N}_E(\rho^V_W) \subseteq \mathcal{N}_E(\rho^{\overline{V}}_{\overline{W}})$ which lies on any geodesic λ connecting $\pi_W(x)$ to $\pi_w(y)$. Connect the endpoints of λ to the endpoints of $\overline{\gamma}$ with two geodesics, each of whose length is at most $\mathsf{diam}(\pi_{\overline{W}}(\overline{x})) \leqslant \beth$ because $\pi_W(x) \subseteq \pi_{\overline{W}}(\overline{x})$ and similarly for y. Since geodesic quadrangles in $\mathcal{C}W$ are $2E$-slim, we obtain $\mathsf{d}_W(w, \overline{\gamma}) \leqslant 2E + \beth$, so $\overline{\gamma} \cap N_{3E+\beth}(\rho^{\overline{V}}_{\overline{W}}) \neq \varnothing$, as desired.

Now suppose $\overline{W} = \overline{S}$. In this case, $\overline{\gamma}$ is a geodesic in $\overline{X}$ from $\overline{x}$ to $\overline{y}$. Complete $\overline{\gamma}$ to a geodesic triangle $\overline{T}$ with vertices $\{\overline{x}, \overline{y}, \overline{v_V}\}$, and let T be a lift of $\overline{T}$ to $\widehat{X}$ with vertices $\{x, y, v_V\}$. Let γ be the lift of $\overline{\gamma}$ inside T. Since $\{x, y, V\}$ are pairwise minimal, $\mathsf{d}_V(x, y) \geqslant \mathsf{d}_{\overline{V}}(\overline{x}, \overline{y}) > E$, so Lemma 4.22 yields that $\mathsf{d}_{\widehat{X}}(v_V, \gamma) \leqslant 2$, concluding the proof of the bounded geodesic image axiom.

(10) Partial realization: Let $\{\overline{V}_j\}$ be a collection of pairwise orthogonal domains of $\mathfrak{S}/N$, and fix $p_j \in \mathcal{C}\overline{V}_j$ for each j. If $\{\overline{V}_j\} = \{\overline{S}\}$, then the unique given point is $\overline{p} \in \widehat{X}/N$. Choose any representative $p \in \overline{p}$, and let $x \in X \subseteq \widehat{X}$ be at distance at most one from p. Since vertices of X are all elements of G, we can consider $\overline{x}$ as an element of G/N. Note that $\pi_{\overline{S}}(x) = \overline{x}$, which is at distance 1 from $\overline{p}$ in $\widehat{X}/N$. Hence the first bullet point of the partial realization axiom holds, and the other two vacuously hold.

Thus suppose that $\overline{V}_j \neq \overline{S}$ for all j. Since the $\overline{V}_j$ are pairwise orthogonal, Lemma 4.31 provides a collection $\{V_j\}$ of pairwise orthogonal, pairwise minimal representatives $V_j \in \overline{V}_j$. As $\mathcal{C}\overline{V}_j = \mathcal{C}V_j$ in this case, we have $p_j \in \mathcal{C}V_j$. Let $x \in G$ be the point provided by the partial realization axiom in $(G, \mathfrak{S})$. Since $S \supsetneq V_j$ for each j, we have that $\mathsf{d}_X(x, \rho^S_{V_j}) \leqslant E$. Thus $\{x, V_j\}$ is a minimal pair for each j, since $x \in \mathcal{N}_A(\rho^S_{V_j})$ and therefore x and v_{V_j} are joined by an edge. We will show that $\overline{x} \in G/N$ satisfies the conclusions of the partial realization axiom, Item (10).

The first bullet point of Item (10) holds because $\mathcal{C}\overline{V}_j = \mathcal{C}V_j$ and $\{x, V_j\}$ is minimal for each j, which implies that $\mathsf{d}_{\overline{V}_j}(\overline{x}, p_j) \leqslant \mathsf{d}_{V_j}(x, p_j) \leqslant E$ by Lemma 4.33.

For the second bullet point of the axiom, fix $\overline{W} \in \mathfrak{S}/N$ such that $\overline{V}_j \subsetneq \overline{W}$ or $\overline{V}_j \pitchfork \overline{W}$ for some j, and for simplicity let $\overline{V} = \overline{V}_j$. We will uniformly bound $\mathsf{d}_{\overline{W}}(\overline{x}, \rho^{\overline{V}}_{\overline{W}})$. If $\overline{W} = \overline{S}$ then $\mathsf{d}_{\overline{X}}(\overline{x}, \rho^{\overline{V}}_{\overline{S}}) \leqslant \mathsf{d}_X(x, \rho^V_S) + 1 \leqslant E + 1$ by Lemma 4.33. Otherwise, by Lemma 4.29, there exist $W \in \overline{W}$ and $V' \in \overline{V}$ such that $\{x, V', W\}$ are pairwise minimal. By Lemma 4.20, we can also assume that $V' = V$, as v_V and x are joined by an edge. Hence $\{x, V, W\}$ are pairwise minimal, which in particular means that the relation between V and W is the same as that between $\overline{V}$ and $\overline{V}$. Hence, by minimality and the realization axiom in $\mathfrak{S}$ we have that $\mathsf{d}_{\overline{W}}(\overline{x}, \rho^{\overline{V}}_{\overline{W}}) \leqslant \mathsf{d}_W(x, \rho^V_W) \leqslant E$, concluding the proof.

(11) Uniqueness: Let $\overline{x}, \overline{y} \in G/N$ and $r \geqslant 0$, and suppose that $\mathsf{d}_{\overline{U}}(\overline{x}, \overline{y}) \leqslant r$ for every $\overline{U} \in \mathfrak{S}/N$. We will show that $\mathsf{d}_{G/N}(\overline{x}, \overline{y})$ is uniformly bounded by a constant depending only on r. Let $\{x, y\}$ be minimal representatives with $x \in \overline{x}$ and $y \in \overline{y}$.

Recall that $\widehat{X}$ is the cone-off of X with respect to the family $\mathcal{H} = \mathcal{Y} \cup \{\mathcal{N}_A(\rho^U_S)\}_{U \subsetneq S}$ of $\max\{K, 3E\}$-quasiconvex subsets. By [Spr18, Proposition 2.27], there exists $\xi = \xi(E, K) \geqslant 1$ and a ξ–quasigeodesic γ from x to y in $\widehat{X}$ such that any de-electrification $\widetilde{\gamma}$ is itself a ξ–quasigeodesic from x to y in X. By construction, we have

$$\widetilde{\gamma} = \sigma_1 * \eta_1 * \cdots * \sigma_k * \eta_k * \sigma_{k+1},$$

where each σ_i is a ξ-quasigeodesic in X contained in γ, while each η_i is a geodesic in X connecting points in some subspace $Y_i \in \mathcal{Y}$ or in the A-neighborhood of some $\rho^{U_i}_S$. Let α be (the image in X of) a λ–hierarchy path in G from x to y, whose existence is guaranteed by Theorem 4.5. By Lemma 2.1, the unparameterized λ-quasigeodesic α and the ξ-quasigeodesic $\widetilde{\gamma}$ are at Hausdorff distance at most Φ, where Φ depends only on λ, ξ, and E, and so ultimately on E and K.

Let Ψ be such that each subgroup H_Y has a Ψ–cobounded action on Y, which exists by Hypothesis 4.12, and fix

$$\daleth > 2\Phi + 2\Psi + 3E + 2A. \tag{11}$$

We will consider only the collection of paths $\eta_{i_1}, \ldots, \eta_{i_m}$ whose endpoints are at distance greater than $\daleth$ in X, that is, such that $\mathsf{d}_X((\eta_{i_j})_-, (\eta_{i_j})_+) > \daleth$. In particular, since $\daleth > 2A + E = \mathsf{diam}(\mathcal{N}_A(\rho^U_S))$, the endpoints $c'_j = (\eta_{i_j})_-$ and $d'_j = (\eta_{i_j})_+$ must belong to some

$Y \in \mathcal{Y}$. Let $c_j, d_j \in Y$ be points in the same H_Y-orbit at distance at most Ψ from c'_j and d'_j, respectively. Let $a_j, b_j \in \alpha$ be points at distance at most Φ from c'_j, d'_j respectively, so that $\mathsf{d}_X(a_j, c_j) \leqslant \Phi + \Psi$ and $\mathsf{d}_X(b_j, d_j) \leqslant \Phi + \Psi$. See Figure 10.

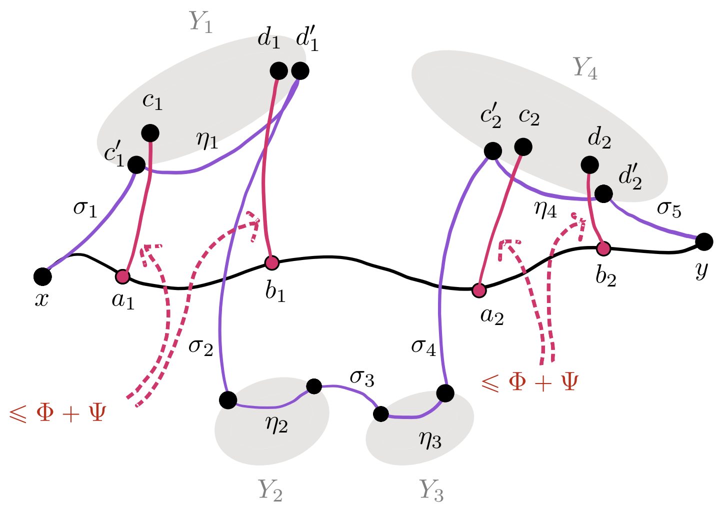

Figure 10: The setup for the proof of the Uniqueness Axiom, in the space X. The de-electrification $\widetilde{\gamma}$ is in purple, while the black path is the λ-hierarchy path α. In this example, $\mathsf{d}_{Y_i}(x, y) \geqslant \daleth$ for $i = 1, 4$.

Let $\omega_1 = \alpha|_{[x,a_1]} \cup [a_1, c_1]$; for $1 < i \leqslant m$, let $\omega_i = [d_{i-1}, b_{i-1}] \cup \alpha|_{[b_{i-1},a_i]} \cup [a_i, c_i]$; and let $\omega_{m+1} = [d_m, b_m] \cup \alpha|_{[b_m,y]}$. See Figure 11.

Recalling that each vertex of X represents an element of G, the path

$$\omega_1 \cup [c_1, d_1] \cup \omega_2 \cup [c_2, d_2] \cup \cdots \cup [c_m, d_m] \cup \omega_{m+1}$$

from x to y shows that we can write the element $x^{-1}y \in G$ as

$$x^{-1}y = (x^{-1}a_1) \cdot (a_1^{-1}c_1) \cdot (c_1^{-1}d_1) \cdot (d_1^{-1}b_1) \cdot \cdots \cdot (b_m^{-1}y).$$

Since c_j, d_j are in the same coset of some H_Y, the element $c_j^{-1}d_j$ belongs to N. In particular, we have

$$x^{-1}y \in w_1N \cdot w_2N \cdot \cdots \cdot w_{m+1}N, \tag{12}$$

where $w_1 = x^{-1}c_1$, and similarly $w_j = d_{j-1}^{-1}c_j$ for $1 < j \leqslant m$ and $w_{m+1} = d_m^{-1}y$.

Recall that the word metric d_G (resp. $\mathsf{d}_{G/N}$) is induced by the finite generating set $\mathcal{T}$ (resp. $\overline{\mathcal{T}}$). In general, if w is a word in the product of cosets $(d_1N)(d_2N)\ldots(d_mN)$, then $|\overline{w}|_{\overline{\mathcal{T}}} \leqslant \sum_{i=1}^{m} |d_i|_{\mathcal{T}}$. Applying this fact to the representation of $x^{-1}y$ in Equation (12), we see that it suffices to bound the $\mathcal{T}$–lengths of the coset representatives w_j. Each is constructed as a product of elements of G, so we bound the $\mathcal{T}$–length of each factor individually.

We proceed via a sequence of claims. The first two show that the $\mathcal{T}$–lengths of the coset representatives is approximated by the lengths of the appropriate paths in X. We first consider factors in coset representatives that lie on α.

Claim 4.35. *There exist constants k_1, k_2 such that if $a, b \in \alpha$, then $\mathsf{d}_G(a, b) \asymp_{k_1,k_2} \mathsf{d}_X(a, b)$.*

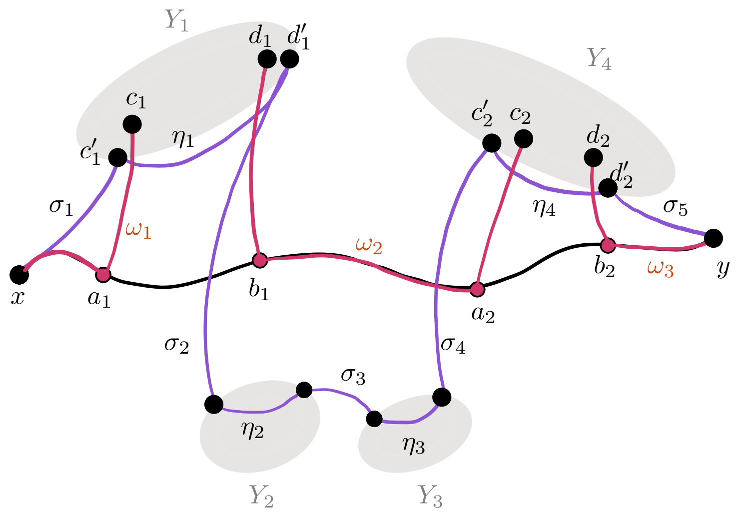

Figure 11: The red segments are the subpaths ω_i in the setup for the proof of the Uniqueness Axiom.

Proof of Claim 4.35. Recall that α is a hierarchy path between x and y, which are minimal. Let $[x, y]$ be an $\widehat{X}$-geodesic. Since x and y are minimal, any two points on $[x, y]$ are minimal, as otherwise we could find closer lifts of $\overline{x}$ and $\overline{y}$. We first prove that $\mathsf{d}_U(x, y)$ is uniformly bounded for all $U \subsetneq S$. Indeed, suppose there exists $U \subsetneq S$ with $\mathsf{d}_U(x, y) > E$. By Lemma 4.22, there exists $z \in [x, y]$ such that $\mathsf{d}_{\widehat{X}}(v_U, z) \leqslant 2$. In turn, since $\{z, x\}$ and $\{z, y\}$ are minimal, Lemma 4.34 provides the existence of $x^* \in \overline{x}$ and $y^* \in \overline{y}$ such that $\{x^*, U\}$ and $\{y^*, U\}$ are minimal and $\max\{\mathsf{d}_U(x, x^*), \mathsf{d}_U(y, y^*)\} \leqslant 9\aleph + 28E$. Hence

$$\mathsf{d}_U(x, y) \leqslant \mathsf{d}_U(x, x^*) + \mathsf{diam}\pi_{\overline{U}}(\overline{x}) + \mathsf{d}_{\overline{U}}(\overline{x}, \overline{y}) + \mathsf{diam}\pi_{\overline{U}}(\overline{y}) + \mathsf{d}_U(y, y^*) \leqslant 18\aleph + 2\beth + 56E + r.$$

Since α is a λ–hierarchy path in G, if $\mathsf{d}_U(x, y)$ is uniformly bounded, then the diameter of $\pi_U(\alpha)$ is at most some uniform constant s'. Therefore $\mathsf{diam}\pi_U(\{a, b\}) \leqslant s'$ whenever $a, b \in \alpha$. Applying the distance formula (Theorem 4.6) with $s = \max\{s' + E + \aleph + 1, s_0\}$ yields constants k_1, k_2 so that $\mathsf{d}_G(a, b) \asymp_{k_1,k_2} \mathsf{d}_X(a, b)$, as required. We note that, for later purposes, we deliberately chose s to be bigger than $\max\{s', s_0\}$, though the latter threshold would have been sufficient to apply the distance formula and conclude the proof of Claim 4.35. □

We next consider factors in coset representatives labeling paths from α to some Y_j.

Claim 4.36. *For $i = 1, \ldots, m$, we have*

$$d_G(a_i, c_i) \asymp_{k_1,k_2} \mathsf{d}_X(a_i, c_i) \leqslant \Phi + \Psi,$$

and similarly

$$d_G(b_i, d_i) \asymp_{k_1,k_2} \mathsf{d}_X(b_i, d_i) \leqslant \Phi + \Psi.$$

Proof of Claim 4.36. We prove the first statement, as the second follows symmetrically. Let $U \in \mathfrak{S} - \{S\}$ be such that $\mathsf{d}_U(a_i, c_i) > E$, so that $[a_i, c_i] \cup \mathcal{N}_E(\rho^U_S) \neq \varnothing$ by the bounded geodesic image axiom. Let $[a_i, c_i] \cup [c_i, d_i] \cup [d_i, b_i] \cup [b_i, a_i]$ be a geodesic rectangle in

X. It follows from the definition of $\daleth$ in Equation (11) that the side $[c_i, d_i]$ has length greater than $2\Phi + 2\Psi + 3E$. Thus $\mathcal{N}_E(\rho_S^U)$ cannot intersect $[b_i, d_i]$ non-trivially, as it already intersects $[a_i, c_i]$. Therefore $\mathsf{d}_U(b_i, d_i) \leqslant E$ by the bounded geodesic image axiom. Moreover, $\mathsf{diam}\pi_U(\{c_i, d_i\}) \leqslant \aleph$ by Remark 4.17, and by the proof of Claim 4.35, we see that $\mathsf{diam}\pi_U(\{a_i, b_i\}) \leqslant s'$, since a_i, b_i lie on α. Thus

$$\mathsf{d}_U(a_i, c_i) \leqslant \mathsf{diam}\pi_U(\{a, b\}) + \mathsf{d}_U(b_i, d_i) + \mathsf{diam}\pi_U(\{c_i, d_i\}) \leqslant s' + E + \aleph.$$

The claim follows by applying the distance formula (Theorem 4.6) with the same threshold $s > s' + E + \aleph$ as in Claim 4.35, and using that $[a_i, c_i]$ and $[b_i, d_i]$ each have X–length at most $\Phi + \Psi$ by construction. □

The final claim relates the sum of the lengths of the subpaths of α to $\mathsf{d}_{\widehat{X}}(x, y)$.

Claim 4.37. *There is a constant $\Xi \geqslant 1$ such that*

$$\mathsf{d}_X(x, a_1) + \sum_{i=1}^{m-1} \mathsf{d}_X(b_i, a_{i+1}) + \mathsf{d}_X(b_m, y) \leqslant \Xi.$$

Proof of Claim 4.37. For convenience, let $b_0 = x$ and $a_{m+1} = y$, so that the sum we are interested in bounding is $\sum_{i=0}^{m} \mathsf{d}_X(b_i, a_{i+1})$. By the triangle inequality, for each $1 \leqslant i \leqslant m-1$ we have that $\mathsf{d}_X(b_i, a_{i+1}) \leqslant \mathsf{d}_X(b_i, d_i') + \mathsf{d}_X(d_i', c_{i+1}') + \mathsf{d}_X(c_{i+1}', a_{i+1})$. Hence

$$\mathsf{d}_X(b_i, a_{i+1}) \leqslant 2\Phi + \ell_X(\widetilde{\gamma}|_{[d_i', c_{i+1}']}) \leqslant 2\Phi + \sum_j \ell_X(\sigma_{i_j}) + \sum_j \ell_X(\eta_{i_j}), \tag{13}$$

where the sums are taken over indices i_j so that each η_{i_j} and σ_{i_j} are contained in the subpath of $\widetilde{\gamma}$ from d_i' to c_{i+1}'. See Figure 12.

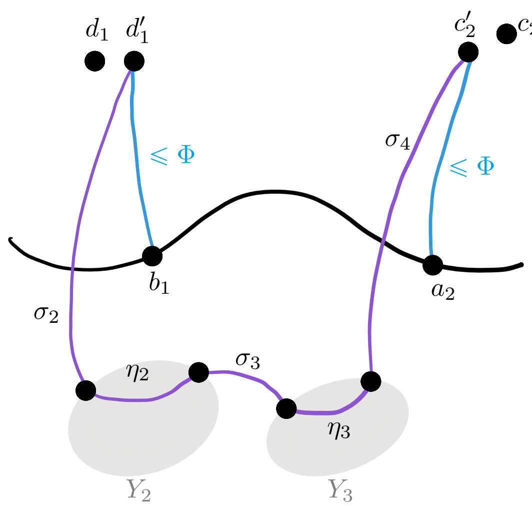

Figure 12: Using the triangle inequality to obtain the bound (13) on $\mathsf{d}_X(b_1, a_2)$ in the proof of Claim 4.37. The purple path is the subpath of the de-electrification $\widetilde{\gamma}$ from d_1' to c_2'.

Similarly

$$\mathsf{d}_X(b_0, a_1) \leqslant \Phi + \sum_j \ell_X(\sigma_{0_j}) + \sum_j \ell_X(\eta_{0_j}) \tag{14}$$

and

$$\mathsf{d}_X(b_m, a_{m+1}) \leqslant \Phi + \sum_j \ell_X(\sigma_{m_j}) + \sum_j \ell_X(\eta_{m_j}), \tag{15}$$

where the sums are taken over indices i_j so that each η_{i_j} and σ_{i_j} are contained in the subpath of $\widetilde{\gamma}$ from x to c_1' and from d_m' to y, respectively.

By our choice of c_i', d_i', each subpath η_{i_j} in any of the sums in (13), (14), and (15) has X–length at most $\daleth$. Moreover, the total number of paths σ_i and η_i is at most the $\widehat{X}$–length of the ξ–quasigeodesic γ, which is in turn bounded by $\xi \mathsf{d}_{\widehat{X}}(x,y)+\xi = \xi(r+1)$. In particular, $m \leqslant \xi(r+1)$, and the sums $\sum_{i=1}^m \sum_j \ell_X(\sigma_{i_j})$ and $\sum_{i=1}^m \sum_j \ell_X(\eta_{i_j})$, where the inner sums are as in (13)–(15), each have at most $\xi(r+1)$ terms. Combining this with the observation that each σ_i is a subpath of the ξ–quasigeodesic γ, we obtain

$$\begin{aligned}
\sum_{i=0}^m \mathsf{d}_X(b_i, a_{i+1}) &\leqslant 2l\Phi + \sum_{i=0}^m \left(\sum_j \ell_X(\sigma_{i_j}) + \sum_j \ell_X(\eta_{i_j}) \right) \\
&\leqslant 2\xi(r+1)\Phi + \sum_{i=0}^m \left(\sum_j \ell_X(\sigma_{i_j}) + \sum_j \daleth \right) \\
&\leqslant \xi(r+1)(2\Phi + \daleth) + \sum_{i=0}^m \sum_j \ell_X(\sigma_{i_j}) \\
&\leqslant \xi(r+1)(2\Phi + \daleth) + \ell_{\widehat{X}}(\gamma) \\
&\leqslant \xi(r+1)(2\Phi + \daleth + 1).
\end{aligned}$$

Thus setting $\Xi = \xi(r+1)(2\Phi + \daleth + 1)$ completes the proof of the claim. □

We are now ready to bound the sum of the lengths of the coset representatives. It follows from Claim 4.35 and Claim 4.36 that for $2 \leqslant i \leqslant m$, we have

$$\begin{aligned}
|w_i|_{\mathcal{T}} = \mathsf{d}_G(d_{i-1}, c_i) &\leqslant \mathsf{d}_G(d_{i-1}, b_{i-1}) + \mathsf{d}_G(b_{i-1}, a_i) + \mathsf{d}_G(a_i, c_i) \\
&\asymp_{k_1,k_2} \mathsf{d}_X(d_{i-1}, b_{i-1}) + \mathsf{d}_X(b_{i-1}, a_i) + \mathsf{d}_X(a_i, c_i) \\
&\preceq_{k_1,k_2} 2\Phi + 2\Psi + \mathsf{d}_X(b_{i-1}, a_i).
\end{aligned}$$

For $i = 1, \ldots, m+1$, a similar argument yields $|w_1|_{\mathcal{T}} = d_G(x, c_1) \preceq_{k_1,k_2} \Phi + \Psi + \mathsf{d}_X(x, a_1)$ and $|w_{m+1}|_{\mathcal{T}} = d_G(d_m, y) \preceq_{k_1,k_2} \Phi + \Psi + \mathsf{d}_X(b_m, y)$. Putting this all together, along with the fact that $m \leqslant \xi(r+1)$ as described in the proof of the previous claim, we have

$$\begin{aligned}
\sum_{i=1}^{m+1} |w_i|_{\mathcal{T}} &\preceq_{k_1,k_2} \mathsf{d}_X(x, a_1) + \Phi + \Psi + \sum_{i=2}^m (2\Phi + 2\Psi + \mathsf{d}_X(b_{i-1}, a_i)) + \mathsf{d}_X(b_m, y) + \Phi + \Psi \\
&\preceq_{k_1,k_2} \xi(r+1)(2\Phi + 2\Psi) + \mathsf{d}_X(x, a_1) + \sum_{i=2}^m \mathsf{d}_X(b_{i-1}, a_i) + \mathsf{d}_X(b_m, y) \\
&\preceq_{k_1,k_2} \xi(r+1)(2\Phi + 2\Psi) + \Xi,
\end{aligned}$$

where the final inequality follows by Claim 4.37. This bound is independent of the choice of x and y, completing the proof of the uniqueness axiom.

(12) Large links: By Lemma 4.2, instead of checking the large link axiom, it suffices to prove that the passing up axiom (12') (Passing up) holds for $(G/N, \mathfrak{S}/N)$, for a suitable choice of the hierarchy constant. To this end we let $t > 0$, and we fix the following constants (see Appendix A for a list of where all involved quantities are defined):

$$\begin{aligned} c_0 &= t + 4\aleph + 20E + 2\beth, \\ c_1 &= 2A + 3E + 2C + J(K + D + 2E) + \Psi, \\ c_2 &= 2c_1 + 2D + 2K, \\ c_3 &= c_0 + c_2 + 12E + 2. \end{aligned}$$

Set $\overline{P} > \max\{P(c_0), (t+1)P(c_3)\}$, where $P\colon \mathbb{R}_{>0} \to \mathbb{N}$ is the function provided by the passing up axiom applied to $(G, \mathfrak{S})$ (with hierarchy constant E).

Let $\overline{V} \in \mathfrak{S}/N$ and $\overline{x}, \overline{y} \in G/N$, and let $\{\overline{U}_1, \overline{U}_2, \ldots, \overline{U}_{\overline{P}}\} \subseteq (\mathfrak{S}/N)_{\overline{V}}$ be such that $\mathsf{d}_{\overline{U}_i}(\overline{x}, \overline{y}) > 5E$ for all $i = 1\ldots, \overline{P}$. Our goal is to find a domain $\overline{W} \sqsubseteq \overline{V}$ and an index i such that $\overline{U_i} \subsetneq \overline{W}$ and $\mathsf{d}_{\overline{W}}(\overline{x}, \overline{y}) > t$.

Case 1: $\overline{V} \neq \overline{S}$. Let $\{x, y, V, U_1, \ldots, U_{\overline{P}}\}$ be pairwise minimal representatives, which exist by Lemma 4.31. Since $\mathsf{d}_{U_i}(x, y) \geqslant \mathsf{d}_{\overline{U}_i}(\overline{x}, \overline{y})$ by Lemma 4.33, the passing up axiom for $(G, \mathfrak{S})$ provides a domain $W \in \mathfrak{S}_V$ containing some U_i such that $\mathsf{d}_W(x, y) > c_0$. Without loss of generality, suppose $U_1 \subsetneq W$. Notice that $\overline{U}_1 \subsetneq \overline{W} \subsetneq \overline{V}$, as pointed out in Remark 4.28. By Lemma 4.29 there exist $x' \in \overline{x}$ and $U_1' \in \overline{U}_i$ such that $\{x', W, U_1'\}$ are pairwise minimal, and we must have that $U_1 = U_1'$ by Proposition 4.25. Then Proposition 4.24 applied to x, x', U, W yields that $\mathsf{d}_W(x, x') \leqslant 2\aleph + 9E$. For the same reason, there exists $y' \in \overline{y}$ with $\{y', W\}$ minimal such that $\mathsf{d}_W(y, y') \leqslant 2\aleph + 9E$. Thus, by Lemma 4.33 we have that

$$\mathsf{d}_{\overline{W}}(\overline{x}, \overline{y}) \geqslant \mathsf{d}_W(x', y') - 2\beth \geqslant \mathsf{d}_W(x, y) - \mathsf{diam}\pi_W(x') - \mathsf{diam}\pi_W(y') - 4\aleph - 18E - 2\beth > t.$$

Case 2: $\overline{V} = \overline{S}$. Towards a contradiction, assume that $\mathsf{d}_{\overline{W}}(\overline{x}, \overline{y}) \leqslant t$ for every $\overline{W}$ which properly contains some $\overline{U}_i$, so that, in particular, $\mathsf{d}_{\overline{X}}(\overline{x}, \overline{y}) \leqslant t$. As in Lemma 4.30, let $\{x, y, U_1, \ldots, U_{\overline{P}}\}$ be such that $\{x, y, U_i\}$ is pairwise minimal for each i. As $\overline{P} > P(c_0)$ by assumption, the passing up axiom for $(G, \mathfrak{S})$ provides a domain $W \in \mathfrak{S}$ properly containing some U_i such that $\mathsf{d}_W(x, y) > c_0$. If $W \neq S$ then, arguing as in Case 1, we obtain $\mathsf{d}_{\overline{W}}(\overline{x}, \overline{y}) > t$, contradicting our assumption. Hence we must have that $\mathsf{d}_X(x, y) > c_0$, while $\mathsf{d}_W(x, y) \leqslant c_0$ for every non-maximal domain W properly containing some U_i. Let $[x, y]$ be an X–geodesic between x and y. By minimality $\mathsf{d}_{U_i}(x, y) \geqslant \mathsf{d}_{\overline{U}_i}(\overline{x}, \overline{y}) > E$, and so the bounded geodesic image axiom (Definition 4.1.(9)) provides a point $p_i \in [x, y] \cap \mathcal{N}_E(\rho_S^{U_i})$ for every $i = 1, \ldots, \overline{P}$. Let $\mathfrak{P}$ be the collection of such points. The set $\mathfrak{P}$ must be "well-distributed" along $[x, y]$, in the following sense.

Claim 4.38. *If a subset $\mathfrak{P}' \subseteq \mathfrak{P}$ has diameter at most c_2, then $|\mathfrak{P}'| \leqslant P(c_3)$.*

Proof of Claim 4.38. If $\min_{z \in \mathfrak{P}'} \mathsf{d}_X(x, z) \leqslant 6E$ let $a = x$. Else, let $a \in [x, y]$ be the vertex such that $\mathsf{d}_X(x, a) = \lfloor \min\{\mathsf{d}_X(x, z) \mid z \in \mathfrak{P}'\} - 6E \rfloor$, where $\lfloor \cdot \rfloor$ denotes the integer part. In particular, $\mathsf{d}(a, \mathfrak{P}) \in [6E, 6E + 1]$. Similarly, if $\min_{z \in \mathfrak{P}'} \mathsf{d}_X(y, z) \leqslant 6E$ let $b = y$; otherwise let $b \in [x, y]$ be the point such that $\mathsf{d}_X(y, b) = \lfloor \min\{\mathsf{d}_X(y, z) \mid z \in \mathfrak{P}'\} - 6E \rfloor$. See Figure 13.

Notice that $\mathsf{d}_{U_i}(x, a) \leqslant E$ for every i such that $p_i \in \mathfrak{P}'$. This is vacuously true if $a = x$; otherwise $\mathsf{d}_X([x, a], \rho_S^{U_i}) > E$, or else the segment between a and $\min \mathfrak{P}'$ would have length at most $3E$. The required inequality now follows from the bounded geodesic image axiom. A symmetric argument holds for y and b, so $\mathsf{d}_{U_i}(a, b) \geqslant \mathsf{d}_{U_i}(x, y) - 4E > E$.

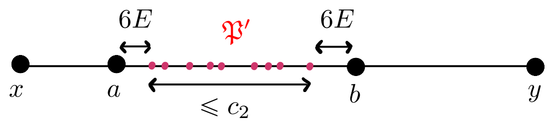

Figure 13: The construction of the points a and b with respect to the subset $\mathfrak{P}'$ (in red) in the proof of Claim 4.38.

If $\mathfrak{P}'$ contained more than $P(c_3)$ points, then the passing up axiom for $(G, \mathfrak{S})$ would imply that $\mathsf{d}_W(a,b) > c_3$ for some W containing some U_i; say $U_1 \subsetneq W$. Notice that W cannot be S, because

$$\mathsf{d}_X(a,b) \leqslant \mathsf{diam}\mathfrak{P}' + 12E + 2 \leqslant c_2 + 12E + 2 \leqslant c_3.$$

Furthermore, if $x \neq a$ then $\mathsf{d}_X([x,a], \rho_S^W) > E$, as otherwise we would have that

$$\mathsf{d}_X(a, \min \mathfrak{P}') \leqslant 2E + \mathsf{diam}(\rho_S^W \cup \rho_S^{U_1}) \leqslant 5E$$

by the consistency axiom Definition 4.1.(7). Thus $\mathsf{d}_W(x,a) \leqslant E$ by the bounded geodesic image axiom, and the same inequality holds for b and y. Hence we again obtain

$$\mathsf{d}_W(x,y) \geqslant \mathsf{d}_W(a,b) - 4E > c_0,$$

contradicting our assumption that x and y are c_0–close in every domain $W \neq S$. □

Now let γ be an $\widehat{X}$–geodesic between x and y, let $\widetilde{\gamma}$ be a de-electrification of γ, and let γ_0 be the vertices of γ lying in X. Note that γ_0 consists of at most $t+1$ vertices, since by minimality $\mathsf{d}_{\widehat{X}}(x,y) = \mathsf{d}_{\overline{X}}(\overline{x}, \overline{y}) \leqslant t$. Let $\mathfrak{Q} \subseteq \mathfrak{P}$ consist of all points in the $(c_2/2)$–neighborhood of γ_0. which therefore contains at most $(t+1)P(c_3)$ points by Claim 4.38. Since $\overline{P} > (t+1)P(c_3)$, there must be some $U' \in \{U_1, \dots, U_{\overline{P}}\}$ such that the corresponding vertex $p' \in \mathfrak{P}$ does not belong to $\mathfrak{Q}$, that is, $\mathsf{d}_X(p', \gamma_0) > c_2/2 = c_1 + D + K$.

Since $p' \in [x,y]^X$, Lemma 2.9 produces a point $q' \in \widetilde{\gamma}$ such that $\mathsf{d}_X(p', q') \leqslant D$. Furthermore $p' \notin \mathfrak{Q}$, which implies that q' lies on a geodesic segment $[c,d]$ of $\widetilde{\gamma} - \gamma_0$ such that $\min\{\mathsf{d}_X(c,q'), \mathsf{d}_X(d,q')\} > c_1 + K$; see Figure 14. Since $c_1 > 2A + E$, the vertices c, d do not belong to $\mathcal{N}_A(\rho_S^V)$ for any $V \in \mathfrak{S}$ such that $v_V \in \gamma$. Hence there exists $Y \in \mathcal{Y}$ such that $v_Y \in \gamma$ and $c, d \in Y$. Let $r' \in \pi_Y(q')$. Since Y is K–quasiconvex, we have $\mathsf{d}_X(r', q') \leqslant K$, and so $\min\{\mathsf{d}_X(c,r'), \mathsf{d}_X(d,r')\} > c_1$.

Using the triangle inequality and that π_Y is J–Lipschitz, we have that

$$\begin{aligned}
\mathsf{d}_Y^\pi(c, v_{U'}) &\geqslant \mathsf{d}_Y^\pi(c, r') - \mathsf{d}_Y^\pi(r', v_{U'}) \\
&\geqslant \mathsf{d}_X(c, r') - J\left(\mathsf{d}_X(r', \rho_S^{U'}) + \mathsf{diam}\rho_S^{U'}\right) \\
&> c_1 - J(K + D + 2E) \geqslant 2C.
\end{aligned}$$

Moreover $\mathsf{d}_Y^\pi(x,c) \leqslant C$, since the geodesic subsegment of γ between x and c does not contain v_Y, so by triangle inequality $\mathsf{d}_Y^\pi(x, v_{U'}) > \mathsf{d}_Y^\pi(c, v_{U'}) - \mathsf{d}_Y^\pi(x,c) > C$. Lemma 4.15 thus yields that every geodesic between x and $v_{U'}$ must pass through v_Y. In particular $\mathsf{d}_{\widehat{X}}(x, v_{U'}) = \mathsf{d}_{\widehat{X}}(x, v_Y) + \mathsf{d}_{\widehat{X}}(v_Y, v_{U'})$, so the path $\eta_x = \gamma|_{[x, v_Y]} * [v_Y, r'] * [r', v_{U'}]$ is a geodesic as it

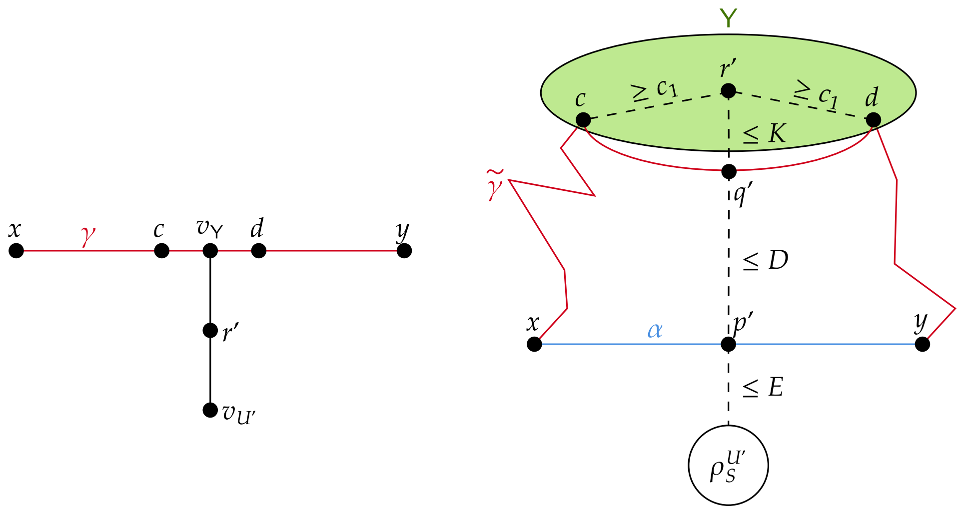

Figure 14: The configuration from the proof of the passing up axiom, as seen from $\widehat{X}$ (on the left) and from X (on the right).

realizes the distance between x and v_U. Similarly, the path $\eta_y = \gamma|_{[y,v_Y]} * [v_Y, r'] * [r', v_{U'}]$ is also a geodesic.

Since H_Y acts Ψ–coboundedly on Y, there exists $h_Y \in H_Y$ such that $d' \coloneqq h_Y d$ is at distance at most Ψ from c. Bend the geodesic tripod with sides $\gamma \cup \eta_x \cup \eta_y$ at v_Y by h_Y. Let η'_y be the image of η_y after bending, and let $y' = h_Y y$. Notice that, since we bent a geodesic triangle, $\{x, y', U'\}$ are again minimal, as so were $\{x, y, U'\}$. The situation is as in Figure 15.

As c is the last point on the geodesic η_x before v_Y, Lemma 4.23 implies that $\mathsf{d}_{U'}(x, c) \leqslant E$. The same argument applied to η'_y yields that $\mathsf{d}_{U'}(d', y') \leqslant E$. Finally, we must have that $\mathsf{d}_{U'}(c, d') \leqslant E$. Indeed, if this was not the case then the bounded geodesic image axiom for $(G, \mathfrak{S})$ would give that $\mathsf{d}_X(\rho_S^{U'}, [c, d']) \leqslant E$ for any X–geodesic $[c, d']$. But this would contradict the fact that $p' \notin \mathfrak{Q}$, as we would have that

$$\mathsf{d}_X(p', c) \leqslant \mathsf{diam}\mathcal{N}_E(\rho_S^{U'}) + \mathsf{d}_X(c, d') \leqslant 3E + \Psi < c_1.$$

Combining the above inequalities, we obtain

$$\mathsf{d}_{U'}(x, y') \leqslant \mathsf{d}_{U'}(x, c) + \mathsf{diam}\pi_{U'}(c') + \mathsf{d}_{U'}(c, d') + \mathsf{diam}\pi_{U'}(d') + \mathsf{d}_{U'}(d', y') \leqslant 5E.$$

However this is a contradiction, because $\{x, y', U'\}$ are pairwise minimal and therefore $\mathsf{d}_{U'}(x, y') \geqslant \mathsf{d}_{\overline{U'}}(\overline{x}, \overline{y}) > 5E$. This concludes the proof of the passing up axiom.

Relative HHG structure: In order to complete the proof of Theorem 4.13, it remains to check that $\mathfrak{S}/N$ is a relative hierarchically hyperbolic *group* structure on G/N. The cofinite G–action on $\mathfrak{S}$ induces a cofinite action of G/N on $\mathfrak{S}/N$. The action preserves the relations $\sqsubseteq$ and $\perp$, since we noticed in Remark 4.28 that two domains in $\mathfrak{S}/N$ are nested (resp. orthogonal) if and only if they admit nested (resp. orthogonal) representatives. Furthermore, the isometries $g\colon \mathcal{C}U \to \mathcal{C}(gU)$ in $(G, \mathfrak{S})$ descend to isometries $\overline{g}\colon \mathcal{C}\overline{U} \to \mathcal{C}(\overline{g}\overline{U})$.

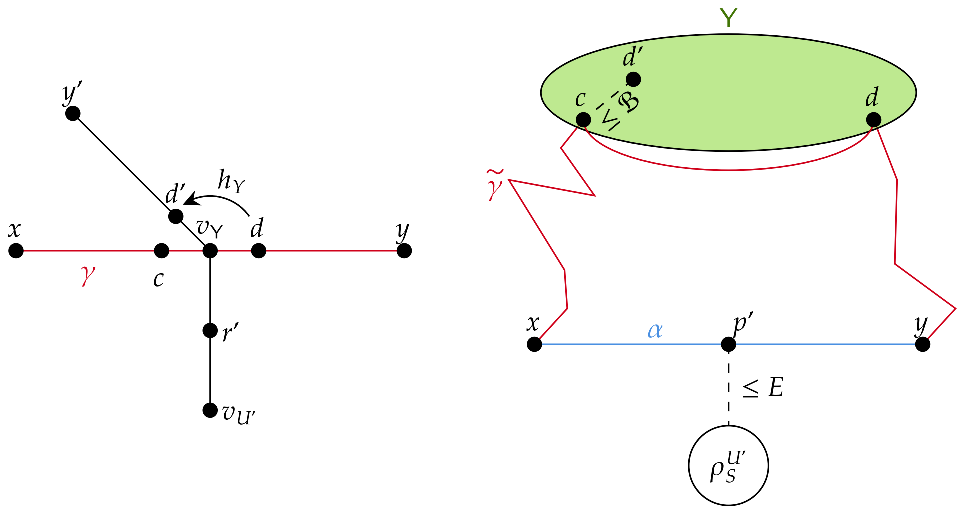

Figure 15: The configuration from the proof of the Passing up axiom, after bending the geodesic connecting y to $v_{U'}$. The points $\{x, y'\}$ are still minimal with U'.

That these isometries satisfy the conditions from Definition 4.7 follows immediately from the fact that they satisfy those conditions in $(G, \mathfrak{S})$, along with the fact that all projections and relative projections between domains are defined in terms of minimal representatives, which are permuted by the G–action on $\widehat{X}$. □

5 Quotients by random walks

This section introduces background on random walks on acylindrically hyperbolic groups and connections with spinning families. These tools will then be used to provide proofs for Theorem B, Corollary C, and Corollary D.

5.1 Background on random walks

Let μ be a probability distribution on a group G. We denote by $\mathrm{Supp}(\mu)$ the *support* of μ, that is, the set of elements $g \in G$ such that $\mu(g) > 0$. Let Γ_μ be the semi-group generated by the support of μ. If Γ_μ is, in fact, a subgroup of G, then μ is called *reversible*. We say μ is *countable* if $\mathrm{Supp}(\mu)$ is countable, is *finitely supported* if $\mathrm{Supp}(\mu)$ is finite, and has *full support* if $\Gamma_\mu = G$. Given a fixed acylindrical action of G on a hyperbolic metric space X, the probability distribution μ is *bounded* if some (equivalently, every) orbit of $\mathrm{Supp}(\mu)$ is a bounded subset of X and *non-elementary* if the action of Γ_μ on X is non-elementary.

Given a reversible, non-elementary probability distribution μ on an acylindrically hyperbolic group G, there exists a unique maximal finite subgroup of G normalized by Γ_μ [Hul16, Lemma 5.5]. We denote this subgroup by $\mathcal{E}_G(\mu)$, or just $\mathcal{E}(\mu)$ when G is understood. We note that $\mathcal{E}(\mu)$ will always contain the maximal finite normal subgroup of G, which we denote by $\mathcal{E}(G)$.

Definition 5.1. The measure μ is *permissible* (with respect to X) if it is bounded, countable, reversible, non-elementary, and $\mathcal{E}(\mu) = \mathcal{E}(G)$.

Remark 5.2. A canonical example of a permissible probability measure is when G is finitely generated and the support of μ is a finite symmetric generating set of G. In this case μ will be finitely supported, hence countable and bounded for any action of G. In addition, such μ will have full support and hence be reversible and non-elementary for any non-elementary, acylindrical action of G. The fact that $\Gamma_\mu = G$ also implies that $\mathcal{E}(\mu) = \mathcal{E}(G)$.

Hypothesis 5.3. Throughout this section, we fix a group G, a cobounded, acylindrical action of G on a δ–hyperbolic space (X, d), and a basepoint $x_0 \in X$. Let $\mu_1, \ldots, \mu_k$ be k permissible probability measures, and let $w_{1,n}, \ldots w_{k,n}$ be independent random walks of length n, starting at x_0, where the step of each $w_{i,n}$ is chosen according to μ_i. When n is fixed, we often suppress the n in the notation for a random walk for simplicity, writing $w_i = w_{i,n}$. If, moreover, a statement applies to all i, we may simply write w for $w_{i,n}$. We say that a property holds *asymptotically almost surely*, or *a.a.s*, if it holds with probability approaching 1 as $n \to \infty$.

Proposition 5.4 ([MT18, Theorem 1.2])**.** *For all i, there exists $\Delta_i = \Delta_i(G, \mu_i) > 0$, called the* drift *of the random walk, such that $\lim_{n\to\infty} \frac{1}{n}\mathsf{d}(x_0, w_{i,n}x_0) = \Delta_i$ almost surely.*

For an isometry g of a hyperbolic space X, the *asymptotic translation length* of g is

$$\tau(g) = \lim_{k\to\infty} \frac{1}{k}\mathsf{d}(x_0, g^k x_0).$$

An element with non-zero asymptotic translation length is a loxodromic isometry, and thus the following proposition implies that $w_{i,n}$ is asymptotically almost surely loxodromic.

Proposition 5.5 ([MT18, Theorem 1.4])**.** *For each i, $\tau(w_{i,n}) > \Delta_i n$ a.a.s.*

Since the G-action on X is acylindrical, [DGO17, Lemma 6.5] implies the existence of a maximal virtually cyclic subgroup $\mathcal{E}(w)$ of G containing w, which consists of all the elements that stabilize any quasi-axis for w up to finite Hausdorff distance. We call $\mathcal{E}(w)$ the *elementary closure* of w. The following proposition states that the elementary closure of a random element is as small as possible:

Proposition 5.6 ([MS19, Proposition 5.1])**.** *For every i, $\mathcal{E}(w_i) = \mathcal{E}(G) \rtimes \langle w_i \rangle$ a.a.s.*

We now construct a $(2, \delta)$–quasi-axis for w_i, which we call α_i. Let $y \in X$ be such that $\mathsf{d}(y, w_i y) \leqslant \inf_{x\in X} \mathsf{d}(x, w_i x) + \delta$. Given any geodesic $[y, w_i y]$, define

$$\alpha_i \coloneqq \bigcup_{r\in\mathbb{Z}} w_i^r [y, w_i y].$$

Notice that α_i is w_i-invariant by construction. Moreover, if the translation length of w_i is sufficiently large with respect to δ (which happens a.a.s. by Proposition 5.5), then α_i is a 7δ-quasiconvex $(2, \delta)$-quasigeodesic, by [Cou16, Corollary 2.7 and Lemma 3.2].

It is also useful to fix a geodesic γ_i in X from x_0 to $w_i x_0$. If we are considering a single random walk w, we denote the quasi-axis α_i by α, and the geodesic γ_i by γ.

An important tool in the study of random walks on hyperbolic spaces with a G–action is matching estimates. We adapt to quasigeodesics the definition of matching from Maher–Sisto [MS19].

Definition 5.7. Let $A, B \geqslant 0$ and $g \in G$. Two quasigeodesics p and q in X have an (A, B, g)*–match* if there are subpaths $p' \subseteq p$ and $q' \subseteq q$ of diameter at least A such that $\mathsf{d}_{Haus}(gp', q') \leqslant B$. If in addition $p = q$, then p has a (A, B, g)*-self-match*; in this case, we say that p has a *disjoint* (A, B, g)*-self-match* if p' and q' are disjoint. We often drop the element g and/or the constants (A, B) when they are not relevant, and simply speak of a match between p and q.

For the rest of the section, let $\Delta = \min_i \Delta_i$ be the minimum drift among all random walks.

Proposition 5.8 ([MT21, Corollary 9.13])**.** *Let w_1, w_2 be independent random walks of length n with respect to permissible probability measures. For any $0 < \varepsilon < 1$ and any $Q \geqslant 0$, γ_1 and γ_2 do not have a $(\varepsilon\Delta n, Q)$–match a.a.s.*

We note that [MT21, Corollary 9.13] is stated for disjoint subpaths of a single random walk, but the same argument, with only the obvious changes, proves our proposition as stated. We next control the matches of overlapping segments of γ.

Proposition 5.9. *Let w be a random walk of length n with respect to a permissible probability measure. For any $0 < \varepsilon < 1$ and any $Q \geqslant 0$, if γ has a $(\varepsilon\Delta n, Q, g)$-self-match then $g \in \mathcal{E}(w)$ a.a.s.*

Notice that $\mathcal{E}(G)$ might act trivially on X, so we cannot forbid self-matches tout-court.

Proof. This is [AH21, Lemma 2.13], which is stated in the case when $\mathcal{E}(G) = \{1\}$ but whose proof runs verbatim in the general case. □

The following proposition can most likely be extracted from [MT21, Section 11], but we provide a proof for clarity and self-containment.

Proposition 5.10. *Let w_1, w_2 be independent random walks of length n, with respect to permissible probability measures. For every $0 < \varepsilon < 1$ and every $Q \geqslant 0$, the following hold a.a.s.:*

- *the axes α_1 and α_2 do not have a $(\varepsilon\Delta n, Q)$–match; and*
- *if α_1 has a $(\varepsilon\Delta n, Q, g)$-self-match, then $g \in \mathcal{E}(w_1)$.*

Proof. Let $\Omega = \Omega(\delta, 7\delta)$ be the constant from Lemma 2.4, and let $\Phi = \Phi(2, \max\{\Omega, 20\delta\}, \delta)$ be the Morse constant provided by Lemma 2.1. Note that Ω and Φ depend only on δ.

Now, assume the following hold for $i = 1, 2$.

(a) $\tau(w_{i,n}) > \Delta n$.

(b) γ_1 and γ_2 do not have a $(\varepsilon'\Delta n, Q + 4\Phi)$–match, and γ_1 does not have a $(\varepsilon'\Delta n, Q + 4\Phi, g)$–self-match unless $g \in \mathcal{E}(w_1)$.

These properties hold a.a.s. by Proposition 5.5 and, respectively, Proposition 5.8 and Proposition 5.9. Now choose $\varepsilon' \in (0, \varepsilon/4)$, and fix n sufficiently large so that the following hold.

(1) $\Delta n > \Omega$.

(2) $\varepsilon\Delta n/4 - 3Q - 4\Phi > \varepsilon'\Delta n$.

The first is possible because Ω does not depend on n, while the second is possible by our choice of $\varepsilon' < \varepsilon/4$ and the fact that Φ does not depend on n.

We shall show that, if α_1 and α_2 have a $(\varepsilon\Delta n, Q)$–match, then γ_1 and γ_2 have a $(\varepsilon'\Delta n, Q+4\Phi)$–match, contradicting (b). With minimal differences in the proof one can also show that, if α_1 has a $(\varepsilon\Delta n, Q, g)$–self-match with $g \notin \mathcal{E}(w_1)$, then γ_1 has an $(\varepsilon'\Delta n, Q+4\Phi, g)$–self-match. Hence we shall only highlight the parts where the arguments diverge.

First, notice that α_i is contained in $\mathcal{N}_{2\Phi}\left(\bigcup_{j\in\mathbb{Z}} w_i^j\gamma_i\right)$. To see this, let $y_i \in \pi_{\alpha_i}(x_0)$ belong to the closest point projection of x_0 onto α_i. Since α_i is invariant under w_i, it follows that $w_iy_i \in \pi_{\alpha_i}(w_ix_0)$; see Figure 16. Furthermore, since α_i is 7δ-quasiconvex and $\mathsf{d}(y, w_iy) \geqslant \tau(w_i) > \Delta n > \Omega$, Lemma 2.4 yields that any nearest point path $[x_0, y_i] \cup [y_i, w_iy_i] \cup [w_iy_i, w_ix_0]$ is a $(1, \Omega)$–quasigeodesic, and therefore lies in the Φ-neighborhood of γ_i. In turn, the $(2, \delta)$-quasigeodesic $\alpha_i|_{[y_i, w_iy_i]}$ is contained in the Φ–neighborhood of $[y_i, w_iy_i]$, and therefore in the 2Φ-neighborhood of γ_i. By applying w_i^j for any j, we see that $\alpha_i|_{[w_i^jy_i, w_i^{j+1}y_i]}$ is contained in the 2Φ–neighborhood of $w_i^j\gamma_i$, as desired.

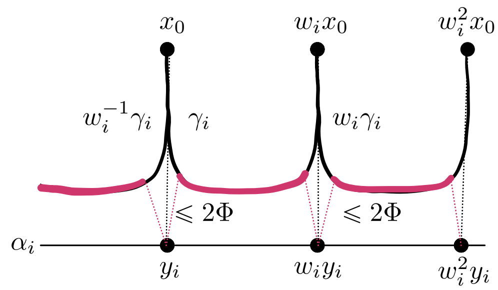

Figure 16: The axis α_i (here, the horizontal line) and the red subpaths of $\bigcup_{j\in\mathbb{Z}} w_i^j\gamma_i$ are at Hausdorff distance at most 2Φ.

Since α_1 and α_2 have an $(\varepsilon\Delta n, Q)$–match, there are subpaths $q \subseteq \alpha_1$ and $p \subseteq \alpha_2$ of diameter $\varepsilon\Delta n$ and $g \in G$ such that $\mathsf{d}_{Haus}(q, gp) \leqslant Q$. In the case of a self-match, $p, q \subseteq \alpha_1$ and $g \notin \mathcal{E}(w_1)$. See Figure 17.

We will consider the orbits $\langle w_1\rangle \cdot y_1$ and $\langle w_2^g\rangle \cdot gy_2$. There are several cases to consider, depending on how these orbits intersect q and gp. Note that w_1 translates points on α_1 by at least $\tau(w_1) \geqslant \Delta n$. Since q has diameter $\varepsilon\Delta n < \Delta n$, at most one orbit point $w_1^jy_1$ can lie on q. Similarly, the intersection between the orbit $\langle w_2^g\rangle \cdot gy_2$ and gp consists of at most one point.

If no orbit point lies on either q or gp, then q and gp are each contained in the 2Φ–neighborhood of some $w_1^{j_1}\gamma_1$ and $gw_2^{j_2}\gamma_2$, respectively. Up to replacing g by $g' = w_1^{-j_1}gw_2^{j_2}$, we can assume that $q \subseteq \mathcal{N}_{2\Phi}(\gamma_1)$ and $p \subseteq \mathcal{N}_{2\Phi}(\gamma_2)$. In the case of a self-match, use $g' = w_1^{-j_1}gw_1^{j_2}$, which again does not belong to $\mathcal{E}(w_1)$. Thus γ_1 and γ_2 have an $(\varepsilon\Delta n - 4\Phi, Q + 4\Phi, g')$–match. By (2), they therefore have an $(\varepsilon'\Delta n, Q+4\Phi, g')$–match, contradicting (b).

Now suppose without loss of generality that an orbit point $w_1^jy_1$ lies on q. Then $w_1^jy_i$ divides q into two subpaths, q_1 and q_2, and one of these has diameter at least $\varepsilon\Delta n/2$. Suppose without loss of generality that it is q_1, and let $z \in gp$ be a point at distance at most Q from $w_1^jy_1$. Consider the subpath gp_1 of gp from its initial point to z. This subpath is at Hausdorff distance at most Q from q_1 and has diameter at least $\varepsilon\Delta n/2 - 2Q$. If no orbit point lies on gp_1, then γ_1 and γ_2 have a $(\varepsilon\Delta n/2 - 2Q - 4\Phi, Q + 4\Phi)$–match; in the case of a self-match, this is realized by some $g' \notin \mathcal{E}(w_1)$. Again by (2), this contradicts (b).

Finally, suppose that an orbit point $gw_2^{j'}y_2$ lies on gp_1. Then this point divides gp_1 into two subpaths, $gp_{1,1}$ and $gp_{1,2}$, one of which has diameter at least $\varepsilon\Delta n/4 - Q$. Without loss of generality, assume it is $gp_{1,1}$; see Figure 17. Let z' be a point on q_1 at distance at most Q from $gw_2^{j'}y_2$. The subpath of q_1 from its initial point to z' is at Hausdorff distance at most Q from $gp_{1,1}$, so it has diameter at least $\varepsilon\Delta n/4 - 3Q$. Again by (2), this contradicts (b).

In all possible configurations of orbit points, we have reached a contradiction with (b), so the proof is complete. □

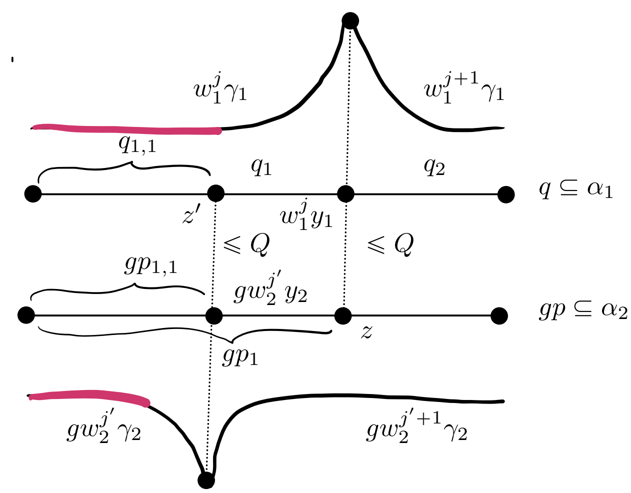

Figure 17: The proof of Proposition 5.10 in the case that a cone point lies on both q and gp. The red subpaths of translates of γ_1 and γ_2 form an $(\varepsilon'\Delta n, Q + 4\Phi)$–match.

5.2 Random walks and spinning families

In this Section we prove Theorem A from the introduction. We first set some notation.

Notation 5.11. Let G be an acylindrically hyperbolic group, so that there exists a cobounded, non-elementary acylindrical action on some hyperbolic space X. Up to replacing X via an equivariant quasi-isometry, we can assume that X is a Cayley graph for G with respect to a (possibly infinite) generating set, so that the action is transitive. Let $\mu_1, \dots, \mu_k$ be permissible probability measures for this action, and consider independent random walks $w_{1,n}, \dots, w_{k,n}$ of length n with respect to these probability measures. For every $i = 1, \dots, k$, let $Y_i = \mathcal{E}(G) \cdot \alpha_i$; notice that $\mathcal{E}(w_i)$ acts on Y_i a.a.s., since $\mathcal{E}(w_i) = \mathcal{E}(G) \rtimes \langle w_i \rangle$ a.a.s by Proposition 5.6. Let $\mathcal{Y}$ be the union of the G-orbits of the Y_i, and for every translate gY_i let $H_{gY_i} = g\langle w_i \rangle g^{-1}$. Finally, let $N = \langle\langle H_{Y_i} \rangle\rangle_{i=1,\dots,k}$. We often drop the indices when they are not relevant.

The following theorem is a more precise version of Theorem A.

Theorem 5.12. *In the setting of Notation 5.11, there are constants E, K, M_0, R, L, where M_0 and L depend on n, such that the collection $(X, \mathcal{Y}, G, \{H_Y\}_{Y\in\mathcal{Y}})$ a.a.s. satisfies Hypothesis 3.1 and the assumption of Corollary 3.16 with respect to (E, K, M_0, R, L). In particular, G/N is a.a.s. acylindrically hyperbolic.*

Proof. Hypothesis 4.12 extends Hypothesis 2.19, so we first verify the latter. First, by assumption X is E-hyperbolic for some $E \geqslant 0$. Moreover, there exists a constant $K = K(E)$ such that every Y is K–quasiconvex. To see this, first notice that, as a consequence of, e.g., [DGO17, Lemma 6.5 & Theorem 6.14], for every $g \in \mathcal{E}(G)$ the translate $g\alpha$ is a $(2, E)$–quasigeodesic with the same ideal endpoints as α. Hence any two $\mathcal{E}(G)$ translates of α lie at Hausdorff distance at most $\Phi = \Phi(2, E, E)$ by Lemma 2.1. By E–slimness of triangles in X, this implies that every geodesic $[x, y]$ with $x, y \in Y$ is in the $(E + \Phi)$–neighborhood of a geodesic between points in the same translate of α. Since α is $7E$–quasiconvex, it follows that Y is K–quasiconvex, where $K = (8E + \Phi)$.

Now fix a constant $0 < \varepsilon < 1$ to be determined later, and let

$$M_0 = \varepsilon\Delta n + 4K + 4E + 2\Phi. \tag{16}$$

We claim that the family $\mathcal{Y}$ is a.a.s. M_0-geometrically separated, as defined in Lemma 2.5. In other words, we have to show that the diameter of $\mathcal{N}_{2K+2E}(Y) \cap Y'$ is at most M_0 for every $Y \neq Y' \in \mathcal{Y}$. Up to translation and relabeling, we may assume $Y = \mathcal{E}(G) \cdot \alpha_1$ and $Y' = g\mathcal{E}(G) \cdot \alpha_i$, where if $i = 1$ then $g \notin \mathcal{E}(w_1)$. Notice that

$$\mathcal{N}_{2K+2E}(Y) \cap Y' \subseteq \mathcal{N}_{2K+2E+\Phi}(\alpha_1) \cap \mathcal{N}_{\Phi}(g\alpha_i).$$

If $\mathsf{diam}(\mathcal{N}_{2K+2E}(Y) \cap Y') \geqslant M_0$, then α_1 and $g\alpha_i$ would have a $(\varepsilon\Delta n, 2K + 2E + 2\Phi, g)$-match, and by Proposition 5.10 this a.a.s. does not occur.

We now prove Hypothesis 3.1. By assumption the action $G \curvearrowright X$ is transitive, i.e., R-cobounded for $R = 0$. Moreover, both $\mathcal{Y}$ and the associated collection $\{H_Y\}_{Y\in\mathcal{Y}}$ of non-trivial subgroups are G-invariant by construction. The next claim shows that we can choose the spinning constant L greater than any given constant which is bounded linearly in M_0. Applying the claim with $\mathcal{L} = \overline{L}$ from Equation (3) will then prove Hypothesis 3.1.(7).

Claim 5.13. *Let $\mathcal{L}(E, K, M_0)$ be a constant which is bounded linearly in terms of M_0. There exist $0 < \varepsilon < 1$ and $L \geqslant 0$ such that the following hold a.a.s.:*

- *$L > \max\{\mathcal{L}(E, K, M_0(\varepsilon)), \varepsilon\Delta n\}$, where M_0 depends on ε as in Equation (16); and*
- *for any $Y \in \mathcal{Y}$, any $x \neq v_Y \in \widehat{X}$, and any $h \in H_Y - \{1\}$, we have $\mathsf{d}^{\pi}_Y(x, hx) > L$.*

Proof of Claim 5.13. Up to the action of G, assume that $Y = Y_i$ for some $i \in \{1, \dots, k\}$, so that $h = w^r_{i,n}$ for some $r \in \mathbb{Z} - \{0\}$. By Proposition 5.5, if $y \in Y_i$ then $\mathsf{d}(y, w^r_{i,n}y) > \Delta n$. As a consequence, for every $x \in \widehat{X} - \{v_{Y_i}\}$ we have that $\mathsf{d}_{Y_i}(x, w^r_{i,n}x) > \Delta n - 2B$, where $B = B(E, K, M_0)$ is the constant from Lemma 2.6 that bounds the diameter of $\pi_{Y_i}(x)$. Since $K = K(E)$, and both B and $\mathcal{L}$ are bounded linearly in M_0, we can find constants $a(E) > 0$, $b(E)$, and $b'(E)$ such that

$$\mathcal{L} + 2B \leqslant a(E)M_0 + b(E) = a(E)\varepsilon\Delta n + b'(E).$$

Now choose ε in such a way that $1 - a(E)\varepsilon > \varepsilon$, and set $L = \Delta n - 2B$. By construction $L \geqslant \max\{\mathcal{L}, \varepsilon\Delta n\}$ for all sufficiently large values of n, as required. $\square$

We now check the assumption of Corollary 3.16. Let w_{k+1} and w_{k+2} be random walks of length n with respect to μ_1 such that $\{w_1, \ldots, w_{k+2}\}$ are pairwise independent. Let α_{k+1} and α_{k+2} be the quasi-axes as in Section 5.1, and let $h, h' \in G$ be such that $x_0 \in h\alpha_{k+1} \cap h'\alpha_{k+2}$. Such elements exist because G acts transitively on X. Finally, let $f = hw_{k+1}h^{-1}$ and $g = h'w_{k+2}(h')^{-1}$, which are a.a.s. loxodromic by Proposition 5.5 and therefore WPD as the action $G \curvearrowright X$ is acylindrical.

We first claim that f and g are a.a.s. independent. If f and g share an ideal endpoint $\xi \in \partial X$, then by Lemma 2.1 the sub-rays $\eta \subseteq h\alpha_{k+1}$ and $\eta' \subseteq h'\alpha_{k+2}$ connecting x_0 to ξ would satisfy $\mathsf{d}_{Haus}(\eta, \eta') \leqslant \Phi$. In particular, α_{k+1} and α_{k+2} would have a $(\Delta n, \Phi)$–match, contradicting Proposition 5.10.

We now show that the axes of f and g are "transverse" to those of the other random walks, in the following sense:

Claim 5.14. $\sup_{Y \in \mathcal{Y}} \sup_{m \in \mathbb{Z}} \mathsf{d}^{\pi}_Y(x_0, f^m x_0) < L/80$ *a.a.s., and similarly for* g.

Proof of Claim 5.14. Suppose toward a contradiction that there exist $n \in \mathbb{Z}$, $i \in \{1, \ldots, k\}$, and $g \in G$ such that $\mathsf{d}^{\pi}_{gY_i}(x_0, f^m x_0) = \mathsf{diam}\pi_{gY_i}(\{x_0, f^m x_0\}) \geqslant L/80$. Notice that $\mathsf{diam}\pi_{gY_i}(x_0)$ does not depend on n, but only on the (uniform) constants E and K, while L grows linearly in n. In particular, if n is sufficiently large it must be the case that $m \neq 0$. Let $y \in \pi_{gY_i}(x_0)$ and $y' \in \pi_{gY_i}(f^m x_0)$ be such that $\mathsf{d}(y, y') \geqslant L/80 - 1$. By Lemma 2.4, if n is sufficiently large then the nearest point path $[x_0, y] \cup [y, y'] \cup [y', f^m x_0]$ is a $(1, \Omega)$-quasigeodesic, where Ω only depends on E. In turn, since x_0 and $f^m x_0$ belong to $h\alpha_{k+1}$, Lemma 2.1 yields that $y, y' \in \mathcal{N}_{\Phi'}(h\alpha_{k+1})$, where $\Phi' = \Phi(2, \Omega, E)$. Let $y'' \in Y$ belong to the same translate of α_i as y, chosen in such a way that $\mathsf{d}(y', y'') \leqslant \Phi$. Then $\mathsf{d}(y, y'') \geqslant L/80 - 1 - \Phi$. Setting $\Phi'' = \max\{\Phi, \Phi'\}$, we have $y, y'' \in \mathcal{N}_{2\Phi''}(h\alpha_{k+1})$. Therefore α_{k+1} and α_i have a $(L/80 - 5\Phi'' - 1, 2\Phi'')$–match. Since $L > \varepsilon\Delta n$, there exists $0 < \chi < 1$ such that $L/80 - 5\Phi'' - 1 > \chi\Delta n$ for large enough values of n. In particular, α_{k+1} and α_i have a $(\chi\Delta n, 2\Phi'')$–match, contradicting Proposition 5.10. ☐

As a consequence of Claim 5.14, f and g are a.a.s. independent loxodromic WPD elements satisfying

$$\sup_{Y \in \mathcal{Y},\, l, m \in \mathbb{Z}} \mathsf{d}^{\pi}_Y(f^m x_0, g^l x_0) \leqslant \sup_{Y \in \mathcal{Y},\, m \in \mathbb{Z}} (\mathsf{d}^{\pi}_Y(x_0, f^m x_0) + \mathsf{d}^{\pi}_Y(x_0, g^m x_0)) < L/40.$$

This shows that the requirements of Corollary 3.16 are a.a.s. satisfied, as required. ☐

We can specialize Theorem 5.12 to the case of relative HHG:

Proposition 5.15. *In the setting of Notation 5.11, assume further that G is an acylindrically hyperbolic relative HHG, and that X is its top-level coordinate space. Then the collection* $(X, \mathcal{Y}, G, \{H_Y\}_{Y \in \mathcal{Y}})$ *a.a.s. satisfies Hypothesis 4.12 with respect to some constants* E, K, M_0, R, L, *where* M_0 *and* L *depend on* n.

Proof. Since Hypothesis 2.19 and Hypothesis 3.1 hold by Theorem 5.12, we are left to check the remaining assumptions of Hypothesis 4.12. It is clear that G acts cofinitely on $\mathcal{Y}$. Furthermore, $\langle w \rangle$ acts geometrically on its axis α, and therefore on Y, since all $\mathcal{E}(G)$-translates of α are within finite Hausdorff distance. Finally, if in Claim 5.13 we choose $\mathcal{L} = \widetilde{L}$ from Equation (6), we can also ensure that $L > \widetilde{L}$, completing the proof of Proposition 5.15. ☐

Remark 5.16. By Corollary 3.9, for every $x \in X$ and every $n \in N - \{1\}$ we have that $\mathsf{d}_X(x, nx) > \tau = (L/10 - 2(B + JR))/J = (L/10 - 2B)/J$; here we used that $R = 0$ as X is a Cayley graph for G. Since B is bounded linearly in terms of M_0, we can choose the constant $\mathcal{L}$ from Claim 5.13 to be larger than $20B$ and ensure that $L/10$ grows faster than $2B$ as $n \to \infty$. This proves that $\tau \to \infty$ as $n \to \infty$, i.e., the subgroup N will a.a.s. have minimum translation length greater than any given constant on the Cayley graph model for X, and therefore on every graph G–equivariantly quasi-isometric to X. This is crucial in the application to random quotients of mapping class groups [Man23].

Remark 5.17. In this paper, we have assumed that all of our random walks have the same length. Some assumptions on the relative lengths of the random walk is necessary for our methods to hold, as we now explain. The matching estimates introduced in Section 5 are key to the arguments in this section. If, for example, both w_1 and w_2 are driven by a uniform measure on the same finite generating set, and the length of w_{1,n_1} is logarithmic in the length of w_{2,n_2}, then w_{1,n_1} will a.a.s. appear as a subword of w_{2,n_2} [ST19, Section 4]. In particular, it will no longer hold that the axes of these two random walks do not have an $(\varepsilon\Delta n_1, Q)$–match a.a.s., as in Proposition 5.10, and so the collection of random walks will not a.a.s. satisfy Hypothesis 4.12. On the other hand, a straightforward generalization of the techniques in this Section should show that Theorem 5.12 holds if the lengths of the random walks differ by linear functions. With more work, it may be possible to extend our methods in the case that n_2 is only *polynomial* in n_1. For simplicity and in the interest of space, we chose not to pursue these directions here.

5.3 Random quotients

We now apply Theorem 5.12 and Proposition 5.15 to prove Theorem B and its corollaries.

Proof of Theorem B. Let $(G, \mathfrak{S})$ be an acylindrically hyperbolic (relative) HHG, and let $\mu_1, \ldots, \mu_k$ be permissible probability measures on G. Let $w_{1,n}, \ldots, w_{k,n}$ be independent random walks of length n with respect to $\mu_1, \ldots, \mu_k$. Theorem 5.12 shows that the G/N-action on $\mathcal{C}S/N$ is non-elementary, and Proposition 5.15 proves that the assumptions of Theorem 4.13 are satisfied. Therefore $(G/N, \mathfrak{S}/N)$ is an acylindrically hyperbolic HHG, as required. □

Proof of Corollary C. Let G be a non-elementary hyperbolic group, and let $w_{1,n}, \ldots, w_{k,n}$ be independent random walks of length n with respect to permissible probability measures on G. Then $(G, \mathfrak{S})$ is a HHG, where $\mathfrak{S} = \{S\}$, and $\mathcal{C}S$ is the Cayley graph of G with respect to a finite generating set. By Theorem B, if $N = \langle\langle w_{1,n}, \ldots, w_{k,n} \rangle\rangle$, then G/N is an acylindrically hyperbolic HHG. Moreover, the hierarchy structure on G/N is $\mathfrak{S}/N = \{\overline{S}\}$, as can be seen from Construction 4.26 in the proof of Theorem 4.13. In particular, $(G/N, \mathfrak{S}/N)$ has no orthogonality, and so is a rank 1 HHG. By [BHS21, Corollary 2.16], G/N is a hyperbolic group, and it is non-elementary as it is acylindrically hyperbolic. □

The proof of Corollary D is similar, but uses relative HHG structures instead of HHG structures.

Proof of Corollary D. Let G be a non-elementary relatively hyperbolic group with infinite, finitely generated peripheral subgroups $\mathcal{H} = \{H_1, \ldots, H_\ell\}$. Let $\mathcal{T}$ be a finite generating set for G such that $\mathcal{T} \cap H$ generates H for every $H \in \mathcal{H}$. By [BHS19, Theorem 9.3], there is a relative HHG structure $(G, \mathfrak{S})$, where:

- $\mathfrak{S} = \{S\} \cup G\mathcal{H}$;
- $\mathcal{C}S = \text{Cay}\left(G, \mathcal{T} \cup \bigcup_{i=1}^{\ell} H_i\right)$, while $\mathcal{C}(gH) = g\text{Cay}(H, \mathcal{T} \cap H)$; and
- every gH is nested in S and transverse to every other domain.

Let $w_{1,n}, \ldots, w_{k,n}$ be independent random walk of length n with respect to permissible probability measures on G, let $N = \langle\langle w_{1,n}, \ldots, w_{k,n} \rangle\rangle$, and let $\overline{\mathcal{T}}$ be the image of $\mathcal{T}$ in G/N. For every $i = 1, \ldots, \ell$, let $\overline{H}_i$ be the image of H_i in G/N, and let $\overline{\mathcal{H}} = \{\overline{H}_1, \ldots, \overline{H}_\ell\}$. Notice that every $H \in \mathcal{H}$ fixes a set of diameter 1 in $\mathcal{C}S$, while every non-trivial $n \in N$ has translation length at least 2 by Corollary 3.9. Hence $H \cap N = \{1\}$, and so $H \cong \overline{H}$.

By Theorem B, G/N is an acylindrically hyperbolic relative HHG with hierarchy structure $\mathfrak{S}/N$. The relation between any two domains of $\mathfrak{S}/N$ is the relation between any minimal representatives, as can be seen from Construction 4.26 in the proof of Theorem 4.13. In particular, $(G/N, \mathfrak{S}/N)$ has no orthogonality, and so is a rank 1 relative HHG. By combining [Rus20, Theorem 4.3] and [Dru09, Proposition 5.1], we conclude that G/N is non-elementarily hyperbolic relative to a collection $\mathcal{Q}$ of subgroups with the property that each $Q \in \mathcal{Q}$ is at finite Hausdorff distance in $\text{Cay}(G/N, \overline{\mathcal{T}})$ from a unique coset of some $\overline{H} \in \overline{\mathcal{H}}$. More precisely, the proof of [Dru09, Proposition 5.1, Step (3)] shows that one such Q exists for every infinite $\overline{H}$, and in particular for every $\overline{H} \in \overline{\mathcal{H}}$ by our assumptions.

Let Q and $\overline{H} \in \overline{\mathcal{H}}$ be as above. Up to replacing Q by some conjugate, we can assume that $\mathsf{d}_{Haus}(Q, \overline{H})$ is finite, say, bounded by some $r \geqslant 0$. We will now show that $\overline{H} = Q$, thus completing the proof of Corollary D. Let $K = Q \cap \overline{H}$, and notice that, by [HW09, Lemma 4.5], there exists a constant $r' > 0$ such that

$$Q \subseteq Q \cap \mathcal{N}_r(\overline{H}) \subseteq \mathcal{N}_{r'}(Q \cap \overline{H}) = \mathcal{N}_{r'}(K).$$

Thus, K has finite index in Q. The same argument shows that K has finite index in $\overline{H}$ as well.

Given $\overline{g} \in \overline{H}$, the subgroup $K \cap \overline{g}K\overline{g}^{-1}$ has finite index in K, as it is the intersection of two finite-index subgroups of $\overline{H}$. Hence $K \cap \overline{g}K\overline{g}^{-1}$ has finite index in Q as well, since K has finite index in Q. Thus $Q \cap \overline{g}Q\overline{g}^{-1}$, which contains $K \cap \overline{g}K\overline{g}^{-1}$, also has finite index in Q. Now Q is infinite, as it is commensurable to $\overline{H} \cong H$, and almost malnormal, as it is a peripheral subgroup in a relative hyperbolic structure. Hence, we must have that $\overline{g} \in Q$. As $\overline{g}$ was an arbitrary element of $\overline{H}$, this proves that $\overline{H} \leqslant Q$.

For the reverse inclusion, suppose toward a contradiction that there is some $\overline{g} \in Q - \overline{H}$. Then $\overline{H}$ and $\overline{g}\overline{H}$ correspond to transverse $\sqsubseteq$–minimal domains in the relative HHG structure of the quotient. Since $\overline{H}$ and $\overline{g}\overline{H}$ are $\sqsubseteq$–minimal, the relative projection $\rho^{\overline{H}}_{\overline{g}\overline{H}}$ is coarsely the nearest point projection of the product region $\mathbf{P}_{\overline{H}}$ to $\mathbf{P}_{\overline{g}\overline{H}}$, and similarly for the other relative projection; see [BHS17a, Remark 1.16] and [Rus20, Lemma 3.1]. Since H is infinite, $\mathbf{P}_{\overline{H}}$ and $\mathbf{P}_{\overline{g}\overline{H}}$ cannot be within finite Hausdorff distance, else we would contradict that the relative projections are bounded diameter sets. On the other hand, $\mathbf{P}_{\overline{H}}$ and $\mathbf{P}_{\overline{g}\overline{H}}$ coarsely coincide with $\overline{H}$ and $\overline{g}\overline{H}$, respectively, and the latter are within finite Hausdorff distance since $\overline{g} \in Q$. This is a contradiction, and so we conclude that $\overline{H} = Q$, as required. □

A List of constants

All constants depending on M_0 are bounded linearly in M_0.

Name	From	Description
δ	Hypothesis 2.19	Hyperbolicity constant of X
$\Phi(\lambda, c, \delta)$	Lemma 2.1	Morse constant for (λ, c)-quasigeodesics
K	Hypothesis 2.19	Quasiconvexity constant of $Y \in \mathcal{Y}$
$J(K, \delta)$	Lemma 2.3	Lipschitz constant of π_Y
$\Omega(\delta, K)$	Lemma 2.4	nearest point path is $(1, \Omega)$-quasigeodesic
$\widehat{\delta}(\delta, K)$	Lemma 2.8	hyperbolicity constant of $\widehat{X}$
$D(\delta, K)$	Lemma 2.9	$[x, y]^X \subseteq \mathcal{N}_D([x, y]^{\widehat{X}})$
M_0	Hypothesis 2.19	$\forall Y \neq Y' \in \mathcal{Y}$, $\mathsf{diam}(Y \cap \mathcal{N}_{2K+2\delta}(Y')) \leqslant M_0$
$M(t) = M(\delta, K, M_0, t)$	Lemma 2.5	$\mathsf{diam}(Y \cap \mathcal{N}_t(Y')) \leqslant M(t)$
$B(\delta, K, M_0)$	Lemma 2.6	$\mathsf{diam}_Y(U) \leqslant B$
$C(\delta, K, M_0)$	Lemma 2.13	Strong BGI: $\mathsf{d}_Y^\pi(x, y) \geqslant C \Rightarrow v_Y \in [x, y]^{\widehat{X}}$
$\theta(\delta, K, M_0)$	Proposition 2.20	$\mathcal{Y}$ satisfies projection axioms wrt θ
R	Hypothesis 3.1	G-action on X is R-cobounded
$\Theta(\delta, K, M_0, R)$	Corollary 2.22	$\widehat{X}$ satisfies projection axioms wrt Θ
$\widetilde{\Theta}(\delta, K, M_0, R)$	Hypothesis 3.1	projection constant with a G-action
$Ж = 33\widetilde{\Theta}$	Definition 2.16	In $\mathcal{P} = \mathcal{P}_{Ж}(\mathcal{Y})$, $W \in \mathrm{Link}(W')$ iff $\forall Y$ $\mathsf{d}_Y(W, W') \leqslant Ж$
L	Hypothesis 3.1	Spinning: $\mathsf{d}_Y^\pi(x, h_Y x) \geqslant L$ $\forall x \in \widehat{X} - \{v_Y\}$
$L_{hyp}(\widetilde{\Theta})$	Remark 3.2	If $L > L_{hyp}$, $\mathcal{P}/N$ is hyperbolic
$\overline{L}(\delta, K, M_0, R) \geqslant L_{hyp}$	Equation (3)	If $L > \overline{L}$, $\widehat{X}/N$ is hyperbolic
$\tau(\delta, K, M_0, R, L)$	Corollary 3.9	$\mathsf{d}_X(x, nx) \geqslant \tau$ $\forall n \in N - \{1\}$.
E	Definition 4.1	relative HHG constant of $(G, \mathfrak{S})$
$A(K, E)$	Equation (5)	in $\widehat{X}$, $\mathcal{N}_A(\rho_S^U)$ is coned off for every $U \subsetneq S$
$\widetilde{L}(E, K, M_0)$	Equation (6)	if $L > \widetilde{L}$ then G/N is a relative HHG
L'	Equation (7)	Spinning constant in X'
$\aleph$	Remark 4.17	$\forall y \in Y$, $\mathsf{diam}(\pi_U(H_Y \cdot y)) \leqslant \aleph$
$\sqsupset(\aleph, E)$	Proposition 4.32	bound on $\mathsf{diam}\pi_{\overline{U}}(\overline{x})$ and $\mathsf{diam}\rho_{\overline{U}}^{\overline{V}}$
Ψ	Theorem 4.13.(11)	H_Y–action on Y is Ψ-cobounded
Δ	Proposition 5.4	minimal drift of random walks

References

[ABD21] Carolyn Abbott, Jason Behrstock, and Matthew G. Durham. Largest acylindrical actions and stability in hierarchically hyperbolic groups. *Trans. Amer. Math. Soc. Ser. B*, 8:66–104, 2021. With an appendix by Daniel Berlyne and Jacob Russell.

[AH21] Carolyn Abbott and Michael Hull. Random walks and quasi-convexity in acylindrically hyperbolic groups. *J. Topol.*, 14(3):992–1026, 2021.

[BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math., Inst. Hautes Étud. Sci.*, 122:1–64, 2015.

[BBFS19] Mladen Bestvina, Ken Bromberg, Koji Fujiwara, and Alessandro Sisto. Acylindrical actions on projection complexes. *Enseign. Math.*, 65(1-2):1–32, 2019.

[BF02] Mladen Bestvina and Koji Fujiwara. Bounded cohomology of subgroups of mapping class groups. *Geom. Topol.*, 6:69–89, 2002.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BHMS24] Jason Behrstock, Mark Hagen, Alexandre Martin, and Alessandro Sisto. A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups. *J. Topol.*, 17(3):94, 2024. Id/No e12351.

[BHS17a] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. Lond. Math. Soc. (3)*, 114(5):890–926, 2017.

[BHS17b] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.

[BHS19] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299:257–338, 2019.

[BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.*, 170(5):909–996, 2021.

[BR22] Daniel Berlyne and Jacob Russell. Hierarchical hyperbolicity of graph products. *Groups Geom. Dyn.*, 16(2):523–580, 2022.

[CdlH16] Yves Cornulier and Pierre de la Harpe. *Metric geometry of locally compact groups*, volume 25 of *EMS Tracts in Mathematics*. European Mathematical Society (EMS), Zürich, 2016. Winner of the 2016 EMS Monograph Award.

[CDP90] Michel Coornaert, Thomas Delzant, and Athanase Papadopoulos. *Géométrie et théorie des groupes. Les groupes hyperboliques de Gromov. (Geometry and group theory. The hyperbolic groups of Gromov)*, volume 1441 of *Lect. Notes Math.* Berlin etc.: Springer-Verlag, 1990.

[CM22] Matt Clay and Johanna Mangahas. Hyperbolic quotients of projection complexes. *Groups Geom. Dyn.*, 16(1):225–246, 2022.

[CMM21] Matt Clay, Johanna Mangahas, and Dan Margalit. Right-angled Artin groups as normal subgroups of mapping class groups. *Compositio Mathematica*, 157(8):1807–1852, 2021.

[Cou16] Rémi B. Coulon. Partial periodic quotients of groups acting on a hyperbolic space. *Ann. Inst. Fourier*, 66(5):1773–1857, 2016.

[Dah18] François Dahmani. The normal closure of big Dehn twists and plate spinning with rotating families. *Geom. Topol.*, 22(7):4113–4144, 2018.

[Del96] Thomas Delzant. Sous-groupes distingués et quotients des groupes hyperboliques. *Duke Math. J.*, 83(3):661–682, 1996.

[DGO17] F. Dahmani, V. Guirardel, and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.*, 245(1156):v+152, 2017.

[DHS20] Matthew Durham, Mark Hagen, and Alessandro Sisto. Correction to the article "Boundaries and automorphisms of hierarchically hyperbolic spaces". *Geom. Topol.*, 24(2):1051–1073, 2020.

[DHS21] François Dahmani, Mark Hagen, and Alessandro Sisto. Dehn filling Dehn twists. *Proc. R. Soc. Edinb., Sect. A, Math.*, 151(1):28–51, 2021.

[Dru09] Cornelia Druţu. Relatively hyperbolic groups: geometry and quasi-isometric invariance. *Comment. Math. Helv.*, 84(3):503–546, 2009.

[Dur23] Matthew Gentry Durham. Cubulating Infinity in Hierarchically Hyperbolic Spaces. Preprint, arXiv:2308.13689 [math.GR], 2023.

[EMM+25] Eduard Einstein, Suraj Krishna MS, Murphy Kate Montee, Thomas Ng, and Markus Steenbock. Random quotients of free products. *arXiv preprint arXiv:2502.08630*, 2025.

[FFMS25] Francesco Fournier-Facio, Giorgio Mangioni, and Alessandro Sisto. Bounded cohomology, quotient extensions, and hierarchical hyperbolicity. Preprint, arXiv:2505.20462 [math.GR], 2025.

[GM08] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.

[Gre60] Martin Greendlinger. *DEHN'S ALGORITHM FOR THE WORD PROBLEM.* ProQuest LLC, Ann Arbor, MI, 1960. Thesis (Ph.D.)–New York University.

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

[Gro93] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[HMS24] Mark Hagen, Alexandre Martin, and Alessandro Sisto. Extra-large type Artin groups are hierarchically hyperbolic. *Math. Ann.*, 388(1):867–938, 2024.

[HRSS25] Mark Hagen, Jacob Russell, Alessandro Sisto, and Davide Spriano. Equivariant hierarchically hyperbolic structures for 3-manifold groups via quasimorphisms. *Ann. Inst. Fourier (Grenoble)*, 75(2):769–828, 2025.

[Hug22] Sam Hughes. Lattices in a product of trees, hierarchically hyperbolic groups and virtual torsion-freeness. *Bull. Lond. Math. Soc.*, 54(4):1413–1419, 2022.

[Hul16] Michael Hull. Small cancellation in acylindrically hyperbolic groups. *Groups Geom. Dyn.*, 10(4):1077–1119, 2016.

[HV24] Sam Hughes and Motiejus Valiunas. Commensurating HNN-extensions: hierarchical hyperbolicity and biautomaticity. *Comment. Math. Helv.*, 99(2):397–436, 2024.

[HW09] G. Christopher Hruska and Daniel T. Wise. Packing subgroups in relatively hyperbolic groups. *Geom. Topol.*, 13(4):1945–1988, 2009.

[KR14] Ilya Kapovich and Kasra Rafi. On hyperbolicity of free splitting and free factor complexes. *Groups Geom. Dyn.*, 8(2):391–414, 2014.

[Man23] Giorgio Mangioni. Random quotients of mapping class groups are quasi-isometrically rigid. Preprint, arXiv:2312.00701 [math.GR], 2023.

[Man24] Giorgio Mangioni. Short hierarchically hyperbolic groups I: uncountably many coarse median structures. Preprint, arXiv:2410.09232 [math.GR], 2024.

[MM99] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999.

[MM00] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.

[MMS24] Joseph Maher, Hidetoshi Masai, and Saul Schleimer. Quotients of the curve complex. *Groups Geom. Dyn.*, 18(2):379–405, 2024.

[MS19] Joseph Maher and Alessandro Sisto. Random subgroups of acylindrically hyperbolic groups and hyperbolic embeddings. *Int. Math. Res. Not.*, 2019(13):3941–3980, 2019.

[MS24] Giorgio Mangioni and Alessandro Sisto. Short hierarchically hyperbolic groups II: quotients and the Hopf property for Artin groups. Preprint, arXiv:2412.04364 [math.GR], 2024.

[MS25] Giorgio Mangioni and Alessandro Sisto. Rigidity of mapping class groups mod powers of twists. *Proceedings of the Royal Society of Edinburgh: Section A Mathematics*, page 1–71, 2025.

[MT18] Joseph Maher and Giulio Tiozzo. Random walks on weakly hyperbolic groups. *J. Reine Angew. Math.*, 742:187–239, 2018.

[MT21] Joseph Maher and Giulio Tiozzo. Random walks, WPD actions, and the Cremona group. *Proc. Lond. Math. Soc.*, 123(2):153–202, 2021.

[Ol'92] Alexander Ol'shanskii. Almost every group is hyperbolic. *International Journal of Algebra and Computation*, 2(1):1–17, 1992.

[Oll04] Y. Ollivier. Sharp phase transition theorems for hyperbolicity of random groups. *Geom. Funct. Anal.*, 14(3):595–679, 2004.

[Osi07] Denis V. Osin. Peripheral fillings of relatively hyperbolic groups. *Inventiones mathematicae*, 167(2):295–326, 2007.

[Osi16] Denis V. Osin. Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.*, 368(2):851–888, 2016.

[PS23] Harry Petyt and Davide Spriano. Unbounded domains in hierarchically hyperbolic groups. *Groups Geom. Dyn.*, 17(2):479–500, 2023.

[Rus20] Jacob Russell. From hierarchical to relative hyperbolicity. *International Mathematics Research Notices*, 2020.

[Rus21] Jacob Russell. Extensions of multicurve stabilizers are hierarchically hyperbolic. arXiv:2107.14116, 2021.

[Spr18] Davide Spriano. Hyperbolic HHS I: Factor systems and quasi-convex subgroups. *arXiv:1711.10931*, 2018.

[ST19] Alessandro Sisto and Samuel J. Taylor. Largest projections for random walks and shortest curves in random mapping tori. *Math. Res. Lett.*, 26(1):293–321, 2019.