

ASYMPTOTIC DIMENSION AND SMALL-CANCELLATION FOR HIERARCHICALLY HYPERBOLIC SPACES AND GROUPS

JASON BEHRSTOCK, MARK F. HAGEN, AND ALESSANDRO SISTO

ABSTRACT. We prove that all hierarchically hyperbolic groups have finite asymptotic dimension. One application of this result is to obtain the sharpest known bound on the asymptotic dimension of the mapping class group of a finite type surface: improving the bound from exponential to at most quadratic in the complexity of the surface. We also apply the main result to various other hierarchically hyperbolic groups and spaces. We also prove a small-cancellation result namely: if G is a hierarchically hyperbolic group, $H \leq G$ is a suitable hyperbolically embedded subgroup, and $N \triangleleft H$ is “sufficiently deep” in H , then G/\widehat{N} is a relatively hierarchically hyperbolic group. This new class provides many new examples to which our asymptotic dimension bounds apply. Along the way, we prove new results about the structure of HHSs, for example: the associated hyperbolic spaces are always obtained, up to quasi-isometry, by coning off canonical coarse product regions in the original space (generalizing a relation established by Masur–Minsky between the complex of curves of a surface and Teichmüller space).

CONTENTS

Introduction	1
1. Preliminaries	6
2. Factored spaces	11
3. Asymptotic dimension of the \mathcal{CU}	15
4. Asymptotic dimension of ball-preimages under $\psi: \mathcal{X} \rightarrow \widehat{\mathcal{X}}$	19
5. Proof of Theorem A and Corollaries B and C	22
6. Quotients of HHG	24
References	35

INTRODUCTION

Motivated by the observation that a suitable CAT(0) cube complex, equipped with a collection of hyperbolic graphs encoding the relationship between its hyperplanes, has many properties exactly parallel to those of the mapping class group, equipped with the collection of curve graphs of subsurfaces, we introduced the class of hierarchically hyperbolic spaces, abbreviated HHS, as a notion of “coarse nonpositive curvature” which provides a framework for studying these two seemingly disparate classes of spaces/groups.

The class of hierarchically hyperbolic spaces consists of metric spaces whose geometry can be recovered, coarsely, from projections onto a specified collection of hyperbolic metric spaces; the axioms governing these spaces and projections are modelled on the relation between

Behrstock was supported as a Simons Fellow.

Hagen was supported by the EPSRC.

Sisto was supported by the Swiss National Science Foundation project 144373.

the subsurface projections between the curve graph of a surface and the curve graph of its subsurfaces, i.e., the mapping class group of a surface is the archetypal HHS (see [MM99, MM00, BKMM12]). The hyperbolic spaces onto which one projects are partially ordered so that there is a unique maximal element and numerous elements that are minimal in every chain in which they appear. Relaxing the hyperbolicity requirement for these minimal spaces, one obtains the notion of a *relatively hierarchically hyperbolic space*, abbreviated RHHS. These notions are reviewed in Section 1 of this paper and a detailed discussion can be found in [BHS14] and [BHS15].

Asymptotic dimension. The *asymptotic dimension* of a metric space is a well-studied quasi-isometry invariant, introduced by Gromov [Gro96], which provides a coarse version of the topological dimension. Early motivation for studying asymptotic dimension was provided by Yu, who showed that groups with finite asymptotic dimension satisfy both the coarse Baum-Connes and the Novikov conjectures [Yu98]. It is now known that asymptotic dimension provides coarse analogues for many properties of topological dimension, see [BD08] for a recent survey. Using very different techniques a number of groups and spaces have been shown to have finite asymptotic dimension, although good estimates on this dimension have proved difficult in many cases: curve graphs [BF07, BB15], mapping class groups [BBF15], cubulated groups [Wri12], graph manifold groups [Smi10], and groups hyperbolic relative to ones with finite asymptotic dimension [Osi05].

One of our main results is the following very general result, which in addition to covering many new cases, provides a unified proof of finite asymptotic dimension for almost all the cases just mentioned:

Theorem A. *Let \mathcal{X} be a uniformly proper HHS. Then $\text{asdim } \mathcal{X} < \infty$. In particular, any HHG has finite asymptotic dimension.*

Theorem A is proven in Section 5, where we establish the slightly stronger Theorem 5.2, which obtains explicit bounds on the dimension.

In special cases where there is a hierarchical structure with a known bound on the asymptotic dimension of the various \mathcal{CU} , we can obtain fairly effective bounds on this dimension, as we now show in the case of the mapping class group.

Bestvina–Bromberg–Fujiwara established finiteness of the asymptotic dimension of mapping class group [BBF15], but without providing explicit bounds. Bestvina–Bromberg then improved on their prior work to obtain an explicit bound on the asymptotic dimension which is exponential in the complexity of the surface [BB15]. Bestvina–Bromberg conjectured that the asymptotic dimension of the mapping class group is equal to its virtual cohomological dimension (and, in particular, linear in the complexity of the surface).

Here, using the hierarchically hyperbolic structure on the mapping class group of a surface, constructed in [BHS15], Section 5 we show that a careful application of Theorem 5.2 yields the following which improves the sharpest bounds from exponential to quadratic:

Corollary B (Asymptotic dimension of $\mathcal{MCG}(S)$). *Let S be a connected oriented surface of finite type of complexity $\xi(S) \geq 2$. Then $\text{asdim } \mathcal{MCG}(S) \leq 5\xi(S)^2$.*

For convenience, we omit from the statement the case of connected oriented surfaces of finite type and complexity at most one (i.e., S^2 with at most 4 punctures and $S^1 \times S^1$ with at most 1 puncture); this is not omitting any cases of interest since such mapping class groups are either finite, \mathbb{Z} , or virtually free, and hence their asymptotic dimensions are 0 or 1.

The notion of a hierarchically hyperbolic structure plays a central role in establishing the bound in Corollary B. Indeed, our proof of Theorem A relies on the fact that “coning off” an appropriate collection of subspaces of an HHS yields a new HHS of lower complexity, and the

bound we eventually obtain is in terms of a uniform bound on the asymptotic dimensions of the hyperbolic spaces in the HHS structure, which is obtained separately in Corollary 3.3.

To establish this bound for general (relatively) hierarchically hyperbolic spaces, we generalize the “tight geodesics” strategy of Bell–Fujiwara [BF07], who proved that the asymptotic dimension of the curve graph of a surface of finite type is finite. Bell–Fujiwara’s work relies on a finiteness theorem of Bowditch [Bow08], which does not provide an explicit bound on the asymptotic dimension.

In the case of the mapping class group, we do not use such upper bounds. Instead, in the proof of Corollary B, we sidestep Section 3 and the “tight geodesics” method completely and instead use the linear bound (in terms of complexity) on the asymptotic dimension of the curve graph provided by Bestvina–Bromberg in [BB15] when we invoke Theorem 5.2. The structure of the proof of Theorem 5.2 allows this, and depends in an essential way on the notion of a hierarchically hyperbolic space. Interestingly, even though the notion of an HHS generalizes known structures in the mapping class group, the present work can not be employed in the mapping class group case without appeal to the full generality of hierarchically hyperbolic spaces. Roughly, this is because our approach involves coning off certain subspaces of a (relatively) hierarchically hyperbolic space to produce a hierarchically hyperbolic space of lower complexity, enabling induction. Although this procedure keeps us in the category of being an HHS, being a mapping class group is not similarly closed under this coning operation.

In [BBF15], the asymptotic dimension of the mapping class group is shown to be finite as a consequence of the fact that the asymptotic dimension of Teichmüller space $\mathcal{T}(S)$ is finite. Our method gives an improved bound on $\text{asdim}(\mathcal{T}(S))$, where $\mathcal{T}(S)$ is given either the Teichmüller metric or the Weil–Petersson metric:

Corollary C. *Let S be a connected oriented surface of finite type of complexity $\xi(S) \geq 1$. Then $\text{asdim}(\mathcal{T}(S)) \leq 5\xi(S)^2 + \xi(S)$.*

Pre-existing bounds on the asymptotic dimension of the associated collection of hyperbolic spaces can be used to bound the asymptotic dimension of some other HHS without recourse to Corollary 3.3. For example, in [BHS14], we showed that if \mathcal{X} is a CAT(0) cube complex admitting a collection of convex subcomplexes called a *factor system*, then \mathcal{X} admits a hierarchically hyperbolic structure in which the associated hyperbolic spaces are uniformly quasi-isometric to simplicial trees, and thus have asymptotic dimension ≤ 1 . This holds in particular when \mathcal{X} embeds convexly in the universal cover of the Salvetti complex of a right-angled Artin group A_Γ associated to a finite simplicial graph Γ . Theorem 5.2 then provides $\text{asdim } \mathcal{X} \leq \sum_{\ell=0}^{|\Gamma^{(0)}|} K_\ell$, where $K_\ell \leq \ell$ is the maximum size of a clique appearing in a subgraph of Γ with ℓ vertices. This reproves finiteness of the asymptotic dimension for such complexes, as established by Wright [Wri12].

We also recover the following result of Osin [Osi05]:

Corollary D (Asymptotic dimension of relatively hyperbolic groups). *Let the group G be hyperbolic relative to a finite collection \mathcal{P} of peripheral subgroups such that $\text{asdim } P < \infty$ for each $P \in \mathcal{P}$. Then $\text{asdim } G < \infty$.*

Proof. It is easy to verify that (G, \mathfrak{S}) is a relatively hierarchically hyperbolic space, where \mathfrak{S} consists of G together with the set of all conjugates of elements of \mathcal{P} ; each of these conjugates is nested in G and the conjugates are pairwise-transverse (the orthogonality relation is empty); see [BHS15]. The result then follows immediately from Theorem 5.2. \square

Quotients of hierarchically hyperbolic groups. As discussed above, the first examples of hierarchically hyperbolic groups were mapping class groups and many cubical groups

[BHS14]; many further constructions and combination theorems were then provided in [BHS15]. In Section 6 we provide many new examples of (relatively) hierarchically hyperbolic groups, which arise as quotients of hierarchically hyperbolic groups by suitable subgroups, using small-cancellation techniques closely related to the theory developed in [DGO11]. In the aforementioned paper, the authors introduced the notion of *hyperbolically embedded subgroup* of a group and extended the relatively hyperbolic Dehn filling theorem [Osi07, GM08], thereby constructing many interesting quotients of groups such as mapping class groups. In particular, they showed that mapping class groups are SQ-universal, i.e. for every hyperbolic surface S and for every countable group Q there exists a quotient of $\mathcal{MCG}(S)$ containing an isomorphic copy of Q . Roughly speaking, we prove that Dahmani-Guirardel-Osin's construction of quotients preserves (relative) hierarchical hyperbolicity when applied to a (relatively) hierarchically hyperbolic group.

We say that the group H is *hierarchically hyperbolically embedded* if G can be generated by a set \mathcal{T} so that: $\mathcal{T} \cap H$ generates H , and $\text{Cay}(G, \mathcal{T})$ is quasi-isometric to the \sqsubseteq -maximal element of \mathfrak{S} , and H is hyperbolically embedded in (G, \mathcal{T}) in the sense of [DGO11]. Theorem E is a direct consequence of Theorem 6.2 below. The proof of Theorem E relies heavily on the “small-cancellation” methods of [DGO11].

Theorem E. *Let (G, \mathfrak{S}) be an HHG and let $H \hookrightarrow_{\text{hh}} (G, \mathfrak{S})$ be hierarchically hyperbolically embedded. Then there exists a finite set $F \subset H - \{1\}$ such that for all $N \triangleleft H$ with $F \cap N = \emptyset$, the group G/\hat{N} is a relatively hierarchically hyperbolic group. If, in addition, H/N is hyperbolic, then G/\hat{N} is hierarchically hyperbolic.*

Here \hat{N} denotes the normal closure of N in G . We remark that acylindrically hyperbolic groups contain plenty of hyperbolically embedded subgroups, and in particular they contain hyperbolically embedded virtually F_2 subgroups [DGO11, Theorem 6.14]. Moreover, hierarchically hyperbolic groups are “usually” acylindrically hyperbolic, in the sense that any non-elementary hierarchically hyperbolic group G so that $\pi_S(G)$ is unbounded is acylindrically hyperbolic, where S is the \sqsubseteq -maximal $S \in \mathfrak{S}$.

Theorem E provides many new examples of hierarchically hyperbolic groups, and hence, via Theorem A, expands the class of groups known to have finite asymptotic dimension. For example:

Corollary F. *Let S be a surface of finite type and let $f \in \mathcal{MCG}(S)$ be a pseudo-Anosov element. Then there exists N such that $\mathcal{MCG}(S)/\langle\langle f^{kN} \rangle\rangle$ is a hierarchically hyperbolic group for all integers $k \geq 1$ and has asymptotic dimension at most $5\xi(S)^2$.*

Corollary F has an exact analogue in the world of cubical groups.

Corollary G. *Let X be a compact special cube complex with universal cover \tilde{X} and let G act properly and cocompactly on \tilde{X} . Let $g \in G$ be a rank-one element, no nonzero power of which stabilizes a hyperplane. Then there exists N such that for all integers $k \geq 1$, the group $G/\langle\langle g^{kN} \rangle\rangle$ is hierarchically hyperbolic and hence has finite asymptotic dimension.*

Since the proofs of Corollaries F and G are very similar, we only give that of Corollary G.

Proof. As shown in [BHS14, BHS15], G is a hierarchically hyperbolic group, where the top-level associated hyperbolic space is quasi-isometric to the intersection graph of the hyperplane carriers in \tilde{X} . As shown in [Hag13], the given g acts loxodromically on this graph, and thus the maximal elementary subgroup containing g is hierarchically hyperbolically embedded in G , see [APMS15, Corollary 3.11] or [Hul13, Corollary 4.14] (which both refine [DGO11, Theorem 6.8]). Finally, apply Theorem E. \square

Note that G need not be virtually special to satisfy the hypotheses of Corollary G (for example, the non-virtually special examples of Burger-Mozes and Wise [BM00, Wis07] act geometrically on the universal cover of the product of two finite graphs), so the corollary can not be proved via cubical small-cancellation theory or related techniques (see [Wis]) followed by an application of the results of [BHS14]. Stronger versions of Corollary F and Corollary G exist, where one kills more complex subgroups.

Remark 1. After we posted the initial version of this paper, it was shown in [HS16] that any proper CAT(0) cube complex admitting a proper, cocompact group action is a hierarchically hyperbolic space (and the group in question a hierarchically hyperbolic group). Hence the conclusion of Corollary G holds with the “special” hypothesis. This is also true for the above-mentioned asymptotic dimension result.

Factored spaces. In section 2 we give a construction which we call a *factored space*. Roughly, the factored space, $\widehat{\mathcal{X}}$, associated to a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ is obtained by collapsing particularly subsets of \mathcal{X} which are isomorphic to direct products. In Proposition 2.4 we prove that this construction yields a new HHS.

One particularly interesting consequence of this construction is the following corollary which is a special case of Corollary 2.9. This result generalizes [MM99, Theorem 1.2] where it is proven that the curve graph of a surface is quasi-isometric to Teichmüller space after collapsing the thin parts, and, also, [MM99, Theorem 1.3] where the mapping class group is considered and the multicurve stabilizers are collapsed.

Corollary H ($\widehat{\mathcal{X}}$ is QI to $\pi_S(\mathcal{X})$). *Let $(\mathcal{X}, \mathfrak{S})$ be hierarchically hyperbolic and let $S \in \mathfrak{S}$ be the unique \sqsubseteq -maximal element. Then $\widehat{\mathcal{X}}$ is quasi-isometric to $\pi_S(\mathcal{X}) \subseteq \mathcal{CS}$.*

This result implies that if (G, \mathfrak{S}) is a hierarchically hyperbolic group, then \mathcal{CS} is quasi-isometric to the coarse intersection graph of the “standard product regions.”

This result can be interpreted as stating that hierarchically hyperbolic structures always arise from a coarse version of the “factor system” construction used in [BHS14] to endow CAT(0) cube complexes with hierarchically hyperbolic structures.

Structure of the paper. In Section 1, we review basic facts about asymptotic dimension and about hierarchical spaces and groups, including (relatively) hierarchically hyperbolic ones. In Section 2, we introduce a coning construction which shows that the top-level hyperbolic space associated to a hierarchically hyperbolic space is quasi-isometric to the space obtained by coning off the “standard product regions.” This construction, which we use in the inductive proof of Theorem A, is of independent interest and generalizes a construction we had originally included in the first version of [BHS15]. Finiteness of the asymptotic dimension of the hyperbolic spaces associated to a relatively hierarchically hyperbolic space, is proved in Section 3. In Section 4, we prove one of the key propositions needed in the induction argument for Theorem A. In Section 5, we prove Theorem A and its corollaries, and finally we prove Theorem E in Section 6.

Acknowledgments. The authors thank the organizers of the conferences *Manifolds and groups* (Ventotene, September 2015) and the *Théorie géométrique et asymptotique des groupes et applications* (CIRM, September 2015), where our initial discussions took place. The authors thank Carolyn Abbott for helpful feedback. The authors also thank the referee for numerous comments that improved the exposition.

1. PRELIMINARIES

1.1. Background on asymptotic dimension. Let (\mathcal{X}, d) be a metric space. There are several equivalent definitions of the *asymptotic dimension* of \mathcal{X} (see e.g. [BD11] or [BD08] for comprehensive surveys). We say that $\text{asdim } \mathcal{X} \leq n$ if for each $D > 0$, there exist $B \geq 0$ and families $\mathcal{U}_0, \dots, \mathcal{U}_n$ of subsets which form a cover of \mathcal{X} such that:

- (1) for all $i \leq n$ and all $U \in \mathcal{U}_i$, we have $\text{diam}(U) \leq B$;
- (2) for all $i \leq n$ and all $U, U' \in \mathcal{U}_i$, if $U \neq U'$ then $d(U, U') > D$.

A function $f: [0, \infty) \rightarrow [0, \infty)$ such that for each sufficiently large D , there is a cover of \mathcal{X} as above that satisfies part (2) for the given D and satisfies part (1) with $B = f(D)$ is an *n-dimensional control function* for \mathcal{X} .

We say \mathcal{X} has *asymptotic dimension n (with control function f)*, when n is minimal so that $\text{asdim } \mathcal{X} \leq n$ (and f is an n -dimensional control function).

A family of metric spaces, $\{(\mathcal{X}_i, d_i)\}_{i \in I}$, has $\text{asdim } \mathcal{X}_i \leq n$ *uniformly* if for all sufficiently large $D \geq 0$, there exists $B \geq 0$ such that for each $i \in I$, there are sets $\mathcal{U}_0^i, \dots, \mathcal{U}_n^i$ of subsets of \mathcal{X}_i , collectively covering \mathcal{X}_i , so that:

- (1) for all $i \in I$, all $0 \leq j \leq n$, and all $U \in \mathcal{U}_j^i$, we have $\text{diam}(U) \leq B$;
- (2) for all $i \in I$, all $0 \leq j < k \leq n$, and all $U, U' \in \mathcal{U}_j^i$, if $U \neq U'$ then $d(U, U') > D$.

As above, $f: [0, \infty) \rightarrow [0, \infty)$ is an *n-dimensional control function* for $\{\mathcal{X}_i\}$ if for each i , and each sufficiently large D , we can choose the covers above so that if the second condition is satisfied for D , then the first is satisfied with $B = f(D)$.

Equivalently, $\text{asdim } \mathcal{X} \leq n$ if for all $r \geq 0$ there exists a uniformly bounded cover of \mathcal{X} such that any r -ball intersects at most $n + 1$ sets in the cover [BD01], and $\{\mathcal{X}_i\}$ has $\text{asdim } \mathcal{X}_i \leq n$ uniformly if for each r the covers can be chosen to consist of sets bounded independently of i . We will use this formulation in Section 3.

We will require the following theorems of Bell–Dranishnikov:

Theorem 1.1 (Fibration theorem; [BD06]). *Let $\psi: \mathcal{X} \rightarrow \mathcal{Y}$ be a Lipschitz map, with \mathcal{X} a geodesic space and \mathcal{Y} a metric space. Suppose that for each $R > 0$, the collection $\{\psi^{-1}(B(y, R))\}_{y \in \mathcal{Y}}$ has $\text{asdim } \psi^{-1}(B(y, R)) \leq n$ uniformly. Then $\text{asdim } \mathcal{X} \leq \text{asdim } \mathcal{Y} + n$.*

Theorem 1.2 (Union theorem; [BD01]). *Let \mathcal{X} be a metric space and assume that $\mathcal{X} = \bigcup_{i \in I} \mathcal{X}_i$, where $\{\mathcal{X}_i\}_{i \in I}$ satisfies $\text{asdim } \mathcal{X}_i \leq n$ uniformly. Suppose that for each R there exists $Y_R \subset \mathcal{X}$, with $\text{asdim } Y_R \leq n$, such that for all distinct $i, i' \in I$, we have $d(\mathcal{X}_i - Y_R, \mathcal{X}_{i'} - Y_R) \geq R$. Then $\text{asdim } \mathcal{X} \leq n$.*

1.2. Background on hierarchical spaces. We recall our main definition from [BHS15]:

Notation 1.3. In Definition 1.4 below, we use the notation $d_W(-, -)$ to denote distance in a space $\mathcal{C}W$, where W is in an index set \mathfrak{S} . We will follow this convention where it will not cause confusion. However, in Section 6, where there are multiple HHS structures and spaces in play, we generally avoid this abbreviation. Similarly, where it will not cause confusion, we write, e.g. $d_W(x, y)$ to mean $d_W(\pi_W(x), \pi_W(y))$, where $x, y \in \mathcal{X}$, and $W \in \mathfrak{S}$, and $\pi: \mathcal{X} \rightarrow \mathcal{C}W$ is a projection. We emphasise that, throughout the text, e.g. $d_W(x, y)$ and $d_{\mathcal{C}W}(\pi_W(x), \pi_W(y))$ mean the same thing.

Definition 1.4 (Hierarchical space, (relative) hierarchically hyperbolic space). The q -quasigeodesic space (\mathcal{X}, d) is a *hierarchical space* if there exists an index set \mathfrak{S} , and a set $\{\mathcal{C}W : W \in \mathfrak{S}\}$ of geodesic spaces $(\mathcal{C}U, d_U)$, such that the following conditions are satisfied:

- (1) **(Projections.)** There is a set $\{\pi_W: \mathcal{X} \rightarrow 2^{\mathcal{C}W} \mid W \in \mathfrak{S}\}$ of *projections* sending points in \mathcal{X} to sets of diameter bounded by some $\xi \geq 0$ in the various $\mathcal{C}W \in \mathfrak{S}$. Moreover, there exists K so that each π_W is (K, K) -coarsely Lipschitz.

- (2) **(Nesting.)** \mathfrak{S} is equipped with a partial order \sqsubseteq , and either $\mathfrak{S} = \emptyset$ or \mathfrak{S} contains a unique \sqsubseteq -maximal element; when $V \sqsubseteq W$, we say V is *nested* in W . We require that $W \sqsubseteq W$ for all $W \in \mathfrak{S}$. For each $W \in \mathfrak{S}$, we denote by \mathfrak{S}_W the set of $V \in \mathfrak{S}$ such that $V \sqsubseteq W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \not\sqsubseteq W$ there is a specified subset $\rho_W^V \subset CW$ with $\text{diam}_{CW}(\rho_W^V) \leq \xi$. There is also a *projection* $\rho_V^W : CW \rightarrow 2^{CV}$. (The similarity in notation is justified by viewing ρ_V^W as a coarsely constant map $CW \rightarrow 2^{CV}$.)
- (3) **(Orthogonality.)** \mathfrak{S} has a symmetric and anti-reflexive relation called *orthogonality*: we write $V \perp W$ when V, W are orthogonal. Also, whenever $V \sqsubseteq W$ and $W \perp U$, we require that $V \perp U$. Finally, we require that for each $T \in \mathfrak{S}$ and each $U \in \mathfrak{S}_T$ for which $\{V \in \mathfrak{S}_T : V \perp U\} \neq \emptyset$, there exists $W \in \mathfrak{S}_T - \{T\}$, so that whenever $V \perp U$ and $V \sqsubseteq T$, we have $V \sqsubseteq W$. Finally, if $V \perp W$, then V, W are not \sqsubseteq -comparable.
- (4) **(Transversality and consistency.)** If $V, W \in \mathfrak{S}$ are not orthogonal and neither is nested in the other, then we say V, W are *transverse*, denoted $V \pitchfork W$. There exists $\kappa_0 \geq 0$ such that if $V \pitchfork W$, then there are sets $\rho_W^V \subseteq CW$ and $\rho_V^W \subseteq CV$ each of diameter at most ξ and satisfying:

$$\min \{d_W(\pi_W(x), \rho_W^V), d_V(\pi_V(x), \rho_V^W)\} \leq \kappa_0$$

for all $x \in \mathcal{X}$.

For $V, W \in \mathfrak{S}$ satisfying $V \not\sqsubseteq W$ and for all $x \in \mathcal{X}$, we have:

$$\min \{d_W(\pi_W(x), \rho_W^V), \text{diam}_{CV}(\pi_V(x) \cup \rho_V^W(\pi_W(x)))\} \leq \kappa_0.$$

The preceding two inequalities are the *consistency inequalities* for points in \mathcal{X} .

Finally, if $U \not\sqsubseteq V$, then $d_W(\rho_W^U, \rho_W^V) \leq \kappa_0$ whenever $W \in \mathfrak{S}$ satisfies either $V \not\sqsubseteq W$ or $V \pitchfork W$ and $W \not\sqsubseteq U$.

- (5) **(Finite complexity.)** There exists $n \geq 0$, the *complexity* of \mathcal{X} (with respect to \mathfrak{S}), so that any set of pairwise- \sqsubseteq -comparable elements has cardinality at most n .
- (6) **(Large links.)** There exist $\lambda \geq 1$ and $E \geq \max\{\xi, \kappa_0\}$ such that the following holds. Let $W \in \mathfrak{S}$ and let $x, x' \in \mathcal{X}$. Let $N = \lambda d_W(\pi_W(x), \pi_W(x')) + \lambda$. Then there exists $\{T_i\}_{i=1, \dots, [N]} \subseteq \mathfrak{S}_W - \{W\}$ such that for all $T \in \mathfrak{S}_W - \{W\}$, either $T \in \mathfrak{S}_{T_i}$ for some i , or $d_T(\pi_T(x), \pi_T(x')) < E$. Also, $d_W(\pi_W(x), \rho_W^{T_i}) \leq N$ for each i .
- (7) **(Bounded geodesic image.)** For all $W \in \mathfrak{S}$, all $V \in \mathfrak{S}_W - \{W\}$, and all geodesics γ of CW , either $\text{diam}_{CV}(\rho_V^W(\gamma)) \leq E$ or $\gamma \cap N_E(\rho_V^W) \neq \emptyset$.
- (8) **(Partial Realization.)** There exists a constant α with the following property. Let $\{V_j\}$ be a family of pairwise orthogonal elements of \mathfrak{S} , and let $p_j \in \pi_{V_j}(\mathcal{X}) \subseteq CV_j$. Then there exists $x \in \mathcal{X}$ so that:
- $d_{V_j}(x, p_j) \leq \alpha$ for all j ,
 - for each j and each $V \in \mathfrak{S}$ with $V_j \not\sqsubseteq V$, we have $d_V(x, \rho_V^{V_j}) \leq \alpha$, and
 - if $W \pitchfork V_j$ for some j , then $d_W(x, \rho_W^{V_j}) \leq \alpha$.
- (9) **(Uniqueness.)** For each $\kappa \geq 0$, there exists $\theta_u = \theta_u(\kappa)$ such that if $x, y \in \mathcal{X}$ and $d(x, y) \geq \theta_u$, then there exists $V \in \mathfrak{S}$ such that $d_V(x, y) \geq \kappa$.

If there exists $\delta \geq 0$ such that \mathcal{CU} is δ -hyperbolic for all $U \in \mathfrak{S}$, then $(\mathcal{X}, \mathfrak{S})$ is *hierarchically hyperbolic*. If there exists δ so that \mathcal{CU} is δ -hyperbolic for all non- \sqsubseteq -minimal $U \in \mathfrak{S}$, then $(\mathcal{X}, \mathfrak{S})$ is *relatively hierarchically hyperbolic*.

We require the following proposition from [BHS15]:

Proposition 1.5 (ρ -consistency). *There exists κ_1 so that the following holds. Suppose that $U, V, W \in \mathfrak{S}$ satisfy both of the following conditions: $U \not\sqsubseteq V$ or $U \pitchfork V$; and $U \not\sqsubseteq W$ or $U \pitchfork W$. Then, if $V \pitchfork W$, then*

$$\min \{d_W(\rho_W^U, \rho_W^V), d_V(\rho_V^U, \rho_V^W)\} \leq \kappa_1$$

and if $V \sqsubset W$, then

$$\min \{d_W(\rho_W^U, \rho_W^V), \text{diam}_{CV}(\rho_V^U \cup \rho_V^W(\rho_W^U))\} \leq \kappa_1.$$

Notation 1.6. Given a hierarchical space $(\mathcal{X}, \mathfrak{S})$, let E be the maximum of all of the constants appearing in Definition 1.4 and Proposition 1.5. Moreover, if $(\mathcal{X}, \mathfrak{S})$ is δ -(relatively) HHS, then we choose $E \geq \delta$ as well.

Notation 1.7. Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchical space and let $\mathcal{U} \subset \mathfrak{S}$. Given $V \in \mathfrak{S}$, we write $V \perp \mathcal{U}$ to mean $V \perp U$ for all $U \in \mathcal{U}$.

We can now prove the following lemma, analogous to Definition 1.4.(4):

Lemma 1.8. *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchical space and let $W \in \mathfrak{S}$ and let $U, V \in \mathfrak{S}_W - \{W\}$ satisfy $U \perp V$. Then $d_W(\rho_W^U, \rho_W^V) \leq 2E$.*

Proof. Apply partial realization (Definition 1.4.(8)). □

The following lemma, which is [BHS15, Lemma 2.5], is used in Section 3:

Lemma 1.9 (Passing large projections up the \sqsubseteq -lattice). *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchical space. For every $C \geq 0$ there exists N with the following property. Let $V \in \mathfrak{S}$, let $x, y \in \mathcal{X}$, and let $\{S_i\}_{i=1}^N \subseteq \mathfrak{S}_V - \{V\}$ be distinct and satisfy $d_{S_i}(x, y) \geq E$. Then there exists $S \in \mathfrak{S}_V$ and i so that $S_i \sqsubset S$ and $d_S(x, y) \geq C$.*

In this paper, we primarily work with relatively HHS. The main results from [BHS15] that we will require are realization, the distance formula, and the existence of hierarchy paths (Theorems 1.12, 1.13, 1.14 below), whose statements require the following definitions:

Definition 1.10 (Consistent tuple). Let $\kappa \geq 0$ and let $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{CU}$ be a tuple such that for each $U \in \mathfrak{S}$, the U -coordinate b_U has diameter $\leq \kappa$. Then \vec{b} is κ -consistent if for all $V, W \in \mathfrak{S}$, we have

$$\min\{d_V(b_V, \rho_V^W), d_W(b_W, \rho_W^V)\} \leq \kappa$$

whenever $V \pitchfork W$ and

$$\min\{d_W(x, \rho_W^V), \text{diam}_V(b_V \cup \rho_V^W)\} \leq \kappa$$

whenever $V \sqsubset W$.

Definition 1.11 (Hierarchy path). A path $\gamma: I \rightarrow \mathcal{X}$ is a (D, D) -hierarchy path if γ is a (D, D) -quasigeodesic and $\pi_U \circ \gamma$ is an unparameterized (D, D) -quasigeodesic for each $U \in \mathfrak{S}$.

Theorem 1.12 (Realization). *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchical space. Then for each $\kappa \geq 1$, there exists $\theta = \theta(\kappa)$ so that, for any κ -consistent tuple $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{CU}$, there exists $x \in \mathcal{X}$ such that $d_V(x, b_V) \leq \theta$ for all $V \in \mathfrak{S}$.*

Observe that uniqueness (Definition (9)) implies that the realization point for \vec{b} provided by Theorem 1.12 is coarsely unique. The following theorem is Theorem 6.7 in [BHS15], which is proved using the corresponding statement for hierarchically hyperbolic spaces ([BHS15, Theorem 4.5]):

Theorem 1.13 (Distance formula for relatively HHS). *Let $(\mathcal{X}, \mathfrak{S})$ be a relatively hierarchically hyperbolic space. Then there exists s_0 such that for all $s \geq s_0$, there exist C, K so that for all $x, y \in \mathcal{X}$,*

$$d(x, y) \asymp_{K, C} \sum_{U \in \mathfrak{S}} \{\{d_U(x, y)\}\}_s.$$

(The notation $\{\{A\}\}_B$ denotes the quantity which is A if $A \geq B$ and 0 otherwise.)

The following closely-related statement is Theorem 6.8 of [BHS15]:

Theorem 1.14 (Hierarchy paths in relatively HHS). *Let $(\mathcal{X}, \mathfrak{S})$ be a relatively hierarchically hyperbolic space. Then there exists $D \geq 0$ such that for all $x, y \in \mathcal{X}$, there is a (D, D) -hierarchy path in \mathcal{X} joining x, y .*

1.2.1. *Hierarchical quasiconvexity, gates, and standard product regions.* The next definition slightly generalizes Definition 5.1 of [BHS15] (which it was stated for the case of hierarchically hyperbolic spaces):

Definition 1.15 (Hierarchical quasiconvexity in relatively HHS). Let $(\mathcal{X}, \mathfrak{S})$ be a δ -relatively hierarchically hyperbolic space, and let $\mathcal{Y} \subseteq \mathcal{X}$. Then \mathcal{Y} is *hierarchically quasiconvex* if there exists a function $k: [0, \infty) \rightarrow [0, \infty)$ such that:

- for each $U \in \mathfrak{S}$ with $\mathcal{C}U$ a δ -hyperbolic space, the subspace $\pi_U(\mathcal{Y}) \subseteq \mathcal{C}U$ is $k(0)$ -quasiconvex;
- for each (\sqsubseteq -minimal) $U \in \mathfrak{S}$ for which $\mathcal{C}U$ is not δ -hyperbolic, either $\mathcal{C}U = \mathcal{N}_{k(0)}^{\mathcal{C}U}(\pi_U(\mathcal{Y}))$ or $\text{diam}(\pi_U(\mathcal{Y})) \leq k(0)$;
- for all $\kappa \geq 0$ and all κ -consistent tuples \vec{b} for which $b_U \subset \pi_U(\mathcal{Y})$ for all $U \in \mathfrak{S}$, each realization point $x \in \mathcal{X}$ for which $d_U(\pi_U(x), b_U) \leq \theta(\kappa)$ satisfies $d(x, \mathcal{Y}) \leq k(\kappa)$ (where $\theta(\kappa)$ is as in Theorem 1.12).

In this case, we say \mathcal{Y} is *k -hierarchically quasiconvex* and refer to k as a *hierarchical quasiconvexity function* for \mathcal{Y} .

Let $(\mathcal{X}, \mathfrak{S})$ be relatively hierarchically hyperbolic and let $\mathcal{Y} \subseteq \mathcal{X}$ be k -hierarchically quasiconvex. Given $x \in \mathcal{X}$ and $U \in \mathfrak{S}$, let $p_U(x)$ be defined as follows. If U is δ -hyperbolic, then $p_U(x)$ is the coarse projection of $\pi_U(x)$ on $\pi_U(\mathcal{Y})$ (which is defined since $\pi_U(\mathcal{Y})$ is $k(0)$ -quasiconvex). If $\pi_U: \mathcal{Y} \rightarrow \mathcal{C}U$ is $k(0)$ -coarsely surjective, then $p_U(x)$ is the set of all $p \in \pi_U(\mathcal{Y})$ with $d_U(x, p) \leq k(0)$ (which is nonempty). Otherwise, $\pi_U(\mathcal{Y})$ has diameter at most $k(0)$, and we let $p_U(x) = \pi_U(\mathcal{Y})$. The tuple $(p_U(x))_{U \in \mathfrak{S}}$ is easily checked to be $\kappa = \kappa(k(0))$ -consistent, and we apply the realization theorem (Theorem 1.12) and the uniqueness axiom to produce a coarsely well-defined point $\mathfrak{g}_{\mathcal{Y}}(x) \in \mathcal{Y}$ so that $d_U(\mathfrak{g}_{\mathcal{Y}}(x), p_U(x))$ is bounded in terms of k for all U . The (coarsely well-defined) map $\mathfrak{g}_{\mathcal{Y}}: \mathcal{X} \rightarrow \mathcal{Y}$ given by $x \mapsto \mathfrak{g}_{\mathcal{Y}}(x)$ is the *gate map* associated to \mathcal{Y} .

Important examples of hierarchically quasiconvex subspaces of the relatively HHS $(\mathcal{X}, \mathfrak{S})$ are the *standard product regions* defined as follows (see [BHS15, Section 5] for more detail). For each $U \in \mathfrak{S}$, let \mathfrak{S}_U denote the set of $V \in \mathfrak{S}$ with $V \sqsubseteq U$, and let \mathfrak{S}_U^\perp denote the set of $V \in \mathfrak{S}$ such that $V \perp U$, together with some $A_U \in \mathfrak{S}$ such that $V \sqsubseteq A_U$ for all V with $V \perp U$. Then there are uniformly hierarchically quasiconvex subspaces $\mathbf{F}_U, \mathbf{E}_U \subseteq \mathcal{X}$ such that $(\mathbf{F}_U, \mathfrak{S}_U), (\mathbf{E}_U, \mathfrak{S}_U^\perp)$ are relatively hierarchically hyperbolic spaces and the inclusions $\mathbf{F}_U, \mathbf{E}_U \hookrightarrow \mathcal{X}$ extend to a uniform quasi-isometric embedding $\mathbf{F}_U \times \mathbf{E}_U \rightarrow \mathcal{X}$ whose image \mathbf{P}_U is hierarchically quasiconvex. We call \mathbf{P}_U the *standard product region* associated to U and, for each $e \in \mathbf{E}_U$, the image of $\mathbf{F}_U \times \{e\}$ is a *parallel copy of \mathbf{F}_U (in \mathcal{X})*. The relevant defining property of \mathbf{P}_U is: there exists α , depending only on $\mathcal{X}, \mathfrak{S}$ and the output of the realization theorem, so that for all $x \in \mathbf{P}_U$ (and hence each parallel copy of \mathbf{F}_U), we have $d_V(x, \rho_V^U) \leq \alpha$ whenever $U \sqsubset V$ or $U \pitchfork V$. Moreover we can choose α so that, if $U \perp V$, then $\text{diam}(\pi_U(\mathbf{F}_V \times \{e\})) \leq \alpha$ for all $e \in \mathbf{E}_V$.

Remark 1.16 (Gates in standard product regions and their factors). Let $(\mathcal{X}, \mathfrak{S})$ be a relatively HHS and let $U \in \mathfrak{S}$. The gate map $\mathfrak{g}_{\mathbf{P}_U}: \mathcal{X} \rightarrow \mathbf{P}_U$ can be described as follows. For each $x \in \mathcal{X}$ and $V \in \mathfrak{S}$, we have:

- $d_V(\pi_V(\mathfrak{g}_{\mathbf{P}_U}(x)), \rho_V^U) \leq \alpha$ if $V \pitchfork U$ or $U \sqsubset V$;
- $\pi_V(\mathfrak{g}_{\mathbf{P}_U}(x)) = \pi_V(x)$ otherwise.

For each $e \in \mathbf{E}_U$, the gate map $\mathfrak{g}_{\mathbf{F}_U \times \{e\}}: \mathcal{X} \rightarrow \mathbf{F}_U \times \{e\}$ is described by:

- $d_V(\pi_V(\mathbf{g}_{\mathbf{F}_U \times \{e\}}(x)), \rho_V^U) \leq \alpha$ if $V \triangleleft U$ or $U \sqsubset V$;
- $\pi_V(\mathbf{g}_{\mathbf{F}_U \times \{e\}}(x)) = \pi_V(x)$ if $V \sqsubseteq U$;
- $d_V(\pi_V(\mathbf{g}_{\mathbf{F}_U \times \{e\}}(x)), \pi_V(e)) \leq \alpha$ if $V \perp U$.

Likewise, for each $f \in \mathbf{F}_U$, the gate map $\mathbf{g}_{\{f\} \times \mathbf{E}_U}: \mathcal{X} \rightarrow \{f\} \times \mathbf{E}_U$ is described by:

- $d_V(\pi_V(\mathbf{g}_{\{f\} \times \mathbf{E}_U}(x)), \rho_V^U) \leq \alpha$ if $V \triangleleft U$ or $U \sqsubset V$;
- $\pi_V(\mathbf{g}_{\{f\} \times \mathbf{E}_U}(x)) = \pi_V(x)$ if $V \perp U$;
- $d_V(\pi_V(\mathbf{g}_{\{f\} \times \mathbf{E}_U}(x)), \pi_V(f)) \leq \alpha$ if $V \sqsubseteq U$.

Remark 1.17 (Standard product regions in relatively HHS). In [BHS15, Section 5.2], standard product regions are constructed in the context of hierarchically hyperbolic spaces. However, the construction uses only the hierarchical space axioms and realization (Theorem 1.12), so that $\mathbf{F}_U, \mathbf{E}_U, \mathbf{P}_U$ can be constructed in an arbitrary hierarchical space. The way we have defined things, the assertion that these subspaces are hierarchically quasiconvex requires $(\mathcal{X}, \mathfrak{S})$ to be relatively hierarchically hyperbolic. It is easy to see, from the definition, that the explanation of hierarchical quasiconvexity from [BHS15] (for HHS) works in the more general setting of relatively HHS.

Definition 1.18 (Totally orthogonal). Given a hierarchical space $(\mathcal{X}, \mathfrak{S})$, we say that $\mathcal{U} \subset \mathfrak{S}$ is *totally orthogonal* if $U \perp V$ for all distinct $U, V \in \mathcal{U}$.

Recall from [BHS15, Lemma 2.1] that there is a uniform bound, namely the complexity, on the size of totally orthogonal subsets of \mathfrak{S} . Observe that if \mathcal{U} is a totally orthogonal set in the relatively HHS $(\mathcal{X}, \mathfrak{S})$, then $\bigcap_{U \in \mathcal{U}} \mathbf{P}_U$ coarsely contains $\prod_{U \in \mathcal{U}} \mathbf{F}_U$.

1.2.2. *Partially ordering relevant domains.* Given a hierarchical space $(\mathcal{X}, \mathfrak{S})$, a constant $K \geq 0$, and $x, y \in \mathcal{X}$, we say that $U \in \mathfrak{S}$ is *K-relevant (for x, y)* if $d_U(x, y) \geq K$. In Section 2 of [BHS15], it is shown that when $K \geq 100E$, then any set $\mathbf{Rel}_{\max}(x, y, K)$ of pairwise \sqsubseteq -incomparable K -relevant elements of \mathfrak{S} can be partially ordered as follows: if $U, V \in \mathbf{Rel}(x, y, K)$, then $U \leq V$ if $U = V$ or if $U \triangleleft V$ and $d_U(\rho_V^U, y) \leq E$. (This was done in [BHS15] in the context of hierarchically hyperbolic spaces, but the arguments do not use hyperbolicity and thus hold for arbitrary hierarchical spaces.)

1.2.3. *Automorphisms and (relatively) hierarchically hyperbolic groups.* Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchical space. An *automorphism* g of $(\mathcal{X}, \mathfrak{S})$ is a map $g: \mathcal{X} \rightarrow \mathcal{X}$, together with a bijection $g^\diamond: \mathfrak{S} \rightarrow \mathfrak{S}$ and, for each $U \in \mathfrak{S}$, an isometry $g^*(U): \mathcal{C}U \rightarrow \mathcal{C}(g^\diamond(U))$ so that the following diagrams coarsely commute whenever the maps in question are defined (i.e., when U, V are not orthogonal):

$$\begin{array}{ccc} \mathcal{X} & \xrightarrow{g} & \mathcal{X}' \\ \pi_U \downarrow & & \downarrow \pi_{g^\diamond(U)} \\ \mathcal{C}(U) & \xrightarrow{g^*(U)} & \mathcal{C}(g^\diamond(U)) \end{array}$$

and

$$\begin{array}{ccc} \mathcal{C}U & \xrightarrow{g^*(U)} & \mathcal{C}(g^\diamond(U)) \\ \rho_V^U \downarrow & & \downarrow \rho_{g^\diamond(V)}^{g^\diamond(U)} \\ \mathcal{C}V & \xrightarrow{g^*(V)} & \mathcal{C}(g^\diamond(V)) \end{array}$$

The finitely generated group G is *hierarchical* if there is a hierarchical structure (G, \mathfrak{S}) on G (equipped with a word-metric) so that the action of G on itself by left multiplication is an action by HS automorphisms (with the above diagrams uniformly coarsely commuting). If (G, \mathfrak{S}) is a (relatively) hierarchically hyperbolic space, we say that (G, \mathfrak{S}) (or just G) is a *(relatively) hierarchically hyperbolic group [(R)HHG]*.

1.3. Very rotating families. In Section 6, we will make use of the *very rotating families* technology introduced in [DGO11]. All of the notions we need in that section are defined there, and we refer the reader to [DGO11] or [Gui14] for additional background.

2. FACTORED SPACES

Given a hierarchical space $(\mathcal{X}, \mathfrak{S})$, we say $\mathfrak{U} \subseteq \mathfrak{S}$ is *closed under nesting* if for all $U \in \mathfrak{U}$, if $V \in \mathfrak{S} - \mathfrak{U}$, then $V \not\sqsubseteq U$.

Definition 2.1 (Factored space). Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchical space. A *factored space* $\widehat{\mathcal{X}}_{\mathfrak{U}}$ is constructed by defining a new metric \hat{d} on \mathcal{X} depending on a given subset $\mathfrak{U} \subset \mathfrak{S}$ which is closed under nesting. First, for each $U \in \mathfrak{U}$, for each pair $x, y \in \mathcal{X}$ for which there exists $e \in \mathbf{E}_U$ such that $x, y \in \mathbf{F}_U \times \{e\}$, we set $d'(x, y) = \min\{1, d(x, y)\}$. For any pair $x, y \in \mathcal{X}$ for which there does not exist such an e we set $d'(x, y) = d(x, y)$. We now define the distance \hat{d} on $\widehat{\mathcal{X}}_{\mathfrak{U}}$. Given a sequence $x_0, x_1, \dots, x_k \in \widehat{\mathcal{X}}_{\mathfrak{U}}$, define its length to be $\sum_{i=1}^{k-1} d'(x_i, x_{i+1})$. Given $x, x' \in \widehat{\mathcal{X}}_{\mathfrak{U}}$, let $\hat{d}(x, x')$ be the infimum of the lengths of such sequences $x = x_0, \dots, x_k = x'$.

Given a hierarchical space $(\mathcal{X}, \mathfrak{S})$, and a set $\mathfrak{U} \subseteq \mathfrak{S}$ closed under nesting, let $\psi: \mathcal{X} \rightarrow \widehat{\mathcal{X}}_{\mathfrak{U}}$ be the set-theoretic identity map. Observe that:

Proposition 2.2. *The map $\psi: \mathcal{X} \rightarrow \widehat{\mathcal{X}}_{\mathfrak{U}}$ is Lipschitz.*

Proof. This follows from the definition of \hat{d} and the fact that \mathcal{X} is a quasigeodesic space. \square

Definition 2.3 (Hat space). Let \mathfrak{U}_1 denote the set of \sqsubseteq -minimal elements of \mathfrak{S} . The *hat space* $\widehat{\mathcal{X}} = \widehat{\mathcal{X}}_{\mathfrak{U}_1}$ is the factored space associated to the set \mathfrak{U}_1 .

Recall that a δ -relatively HHS is an HS $(\mathcal{X}, \mathfrak{S})$, such that for all $U \in \mathfrak{S}$, either $\mathcal{C}U$ is δ -hyperbolic or $U \in \mathfrak{U}_1$.

Proposition 2.4. *Fix a δ -relatively HHS, $(\mathcal{X}, \mathfrak{S})$, and let $\mathfrak{U} \subset \mathfrak{S}$ be closed under \sqsubseteq and contain each $U \in \mathfrak{U}_1$ for which $\mathcal{C}U$ is not δ -hyperbolic. The space $(\widehat{\mathcal{X}}_{\mathfrak{U}}, \mathfrak{S} - \mathfrak{U})$ is an HHS, where the associated $\mathcal{C}(*), \pi_*, \rho_*^*, \sqsubseteq, \perp, \triangleleft$ are the same as in the original structure.*

Proof. We must verify each of the requirements of Definition 1.4. First observe that by the definition of \hat{d} and the fact that (\mathcal{X}, d) is a quasigeodesic space, $(\widehat{\mathcal{X}}_{\mathfrak{U}}, \hat{d})$ is also a (K, K) -quasigeodesic space for some K .

Projections: By our hypothesis on \mathfrak{U} , we have that $\mathcal{C}U$ is δ -hyperbolic for each $U \in \mathfrak{S} - \mathfrak{U}$, so it remains to check that $(\widehat{\mathcal{X}}, \mathfrak{S} - \mathfrak{U})$ is a hierarchical space. The projections $\pi_U: \widehat{\mathcal{X}}_{\mathfrak{U}} \rightarrow 2^{\mathcal{C}U}$ are as before (more precisely, they are compositions of the original projections $\pi_U: \mathcal{X} \rightarrow \mathcal{C}U$ with the set-theoretic identity $\widehat{\mathcal{X}}_{\mathfrak{U}} \rightarrow \mathcal{X}$, but we will abuse notation and call them π_U).

Fix $U \in \mathfrak{S} - \mathfrak{U}$. By Definition 1.4.(1), there exists K , independent of U , so that π_U is (K, K) -coarsely Lipschitz. Let $x, y \in \widehat{\mathcal{X}}_{\mathfrak{U}}$ and let $x = x_0, \dots, x_{\ell} = y$ be a sequence with $\hat{d}(x, y) \geq \sum_{i=0}^{\ell-1} d'(x_i, x_{i+1}) - 1$. Note that $d_U(x, y) \leq \sum_{i=0}^{\ell-1} d_U(x_i, x_{i+1})$.

Let I_1 be the set of $i \in \{0, \dots, \ell - 1\}$ such that $d'(x_i, x_{i+1}) = d(x_i, x_{i+1})$, let I_2 be the set of i for which x_i, x_{i+1} lie in a common parallel copy of \mathbf{F}_V , where $V \triangleleft U$ or $V \not\sqsubseteq U$ and $V \in \mathfrak{U}$, and let I_3 be the set of i so that x_i, x_{i+1} lie in a common parallel copy of \mathbf{F}_W , where $W \perp U$ and $W \in \mathfrak{U}$. Note that we do not need to consider the case where $W \in \mathfrak{U}$ and $U \sqsubseteq W$, since $W \in \mathfrak{U}$ and \mathfrak{U} is closed under nesting. Then

$$d_U(x, y) \leq \sum_{i \in I_1} [Kd'(x_i, x_{i+1}) + K] + 2\alpha|I_2| + \alpha|I_3|.$$

The third term comes from the fact that, given $i \in I_3$ and $W \perp U$ the associated element of \mathfrak{U} with $x_i, x_{i+1} \in \mathbf{F}_W \times \{e\}$ for some $e \in \mathbf{E}_W$, we have that $\pi_U(\mathbf{F}_W)$ has diameter at most α , so $d_U(x_i, x_{i+1}) \leq \alpha$. Combining the above provides the desired coarse Lipschitz constant C .

Nesting, orthogonality, transversality, finite complexity: The parts of Definition 1.4 that only concern \mathfrak{S} and the relations $\sqsubseteq, \perp, \pitchfork$ continue to hold with \mathfrak{S} replaced by $\mathfrak{S} - \mathfrak{U}$. The complexity of $(\hat{\mathcal{X}}_{\mathfrak{U}}, \mathfrak{S} - \mathfrak{U})$ is obviously bounded by that of $(\mathcal{X}, \mathfrak{S})$. (Note that the fact that \mathfrak{U} is closed under nesting is needed to ensure that for all $W \in \mathfrak{S} - \mathfrak{U}$ and $U \sqsubset W$, there exists $V \sqsubset W$ so that $T \sqsubseteq V$ for each T with $T \sqsubset W$ and $T \perp U$.)

Consistency: Since the projections π_* and relative projections ρ_*^* have not changed, consistency holds for $(\hat{\mathcal{X}}_{\mathfrak{U}}, \mathfrak{S} - \mathfrak{U})$ since it holds for $(\mathcal{X}, \mathfrak{S})$.

Bounded geodesic image and large links: The bounded geodesic image axiom holds for $(\hat{\mathcal{X}}_{\mathfrak{U}}, \mathfrak{S} - \mathfrak{U})$ since it holds for $(\mathcal{X}, \mathfrak{S})$ and is phrased purely in terms of geodesics in the various $\mathcal{C}(\ast)$ and relative projections ρ_*^* . The same applies to the large link axiom.

Partial realization: Since for each $U \in \mathfrak{S}$, we have $\pi_U(\mathcal{X}) = \pi_U(\hat{\mathcal{X}}_{\mathfrak{U}})$, and since we have not changed any of the projections π_* or relative projections ρ_*^* , the partial realization axiom for $(\mathcal{X}, \mathfrak{S})$ implies that for $(\hat{\mathcal{X}}_{\mathfrak{U}}, \mathfrak{S} - \mathfrak{U})$.

Uniqueness: This is done in Lemma 2.8 below. \square

Definition 2.5 (Friendly). For $U, V \in \mathfrak{S}$, we say that U is *friendly* to V if $U \sqsubseteq V$ or $U \perp V$. Note that when U is *not* friendly to V , then ρ_U^V is a uniformly bounded subset of \mathcal{CU} .

In the proof of Lemma 2.8, we will need to “efficiently” jump between product regions P_U, P_V . Heuristically, the pairs of points that are “closest in every \mathcal{CW} ” are of the form p, q for some $p \in \mathfrak{g}_{\mathbf{P}_U}(\mathbf{P}_V)$ and $q = \mathfrak{g}_{\mathbf{P}_V}(p)$, and these are the ones we study in the following lemma. In particular, we are interested in the distance formula terms for such pairs p, q .

Lemma 2.6 (Knowing who your friends are). *Let $U, V \in \mathfrak{S}$ and let $p \in \mathfrak{g}_{\mathbf{P}_U}(\mathbf{P}_V)$. For $q = \mathfrak{g}_{\mathbf{P}_V}(p)$, the following holds. If $W \in \mathfrak{S}$ satisfies $d_W(p, q) \geq 10^3 \alpha E$ then W is not friendly to either of U or V , and $d_W(\rho_W^U, \rho_W^V) \geq 500 \alpha E$.*

Proof. If $W \sqsubseteq V$ or $W \perp V$, then $\pi_W(p), \pi_W(q)$ coarsely coincide by the definition of gates. Hence W is not friendly to V . Suppose now that W is friendly to U . Choose $p_0 \in \mathbf{P}_V$ so that $p = \mathfrak{g}_{\mathbf{P}_U}(p_0)$. Since W is not friendly to V and $p_0, p \in \mathbf{P}_V$, the W -coordinates of p_0, q both coarsely coincide with ρ_W^V . Hence, since W is friendly to U , the W -coordinate of p also coarsely coincides with ρ_W^V , contradicting $d_W(p, q) \geq 10^3 \alpha E$. Hence W is not friendly to U . The final assertion follows from the fact that $d_W(p, \rho_W^U), d_W(q, \rho_W^V) \leq E$. \square

Lemma 2.7. *Suppose that $W \in \mathfrak{S}$ and $x, y \in \mathcal{X}$ satisfy $d_W(x, y) \leq 100\alpha$, while $d_V(x, y) \geq 100\alpha$ for some $V \in \mathfrak{S}$. Suppose that W is not friendly to V . Then $d_W(\rho_W^V, x) \leq 200\alpha E$.*

Proof. First suppose $W \pitchfork V$. The lower bound on $d_V(x, y)$ implies that either $d_V(x, \rho_V^W) > E$ or $d_V(y, \rho_V^W) > E$. In the first case, an application of consistency yields the desired conclusion. In the second case, apply consistency and the upper bound on $d_W(x, y)$.

Next suppose $V \sqsubset W$. If $d_W(\rho_W^V, x) \leq 200\alpha E$, then the bound on $d_W(x, y)$ implies that geodesics from $\pi_W(x)$ to $\pi_W(y)$ would remain E -far from ρ_W^V . But then consistency and bounded geodesic image would imply that $\pi_V(x), \pi_V(y)$ lie $\leq E$ -close, a contradiction. \square

Lemma 2.8 (Uniqueness). *For all $\kappa \geq 0$, there exists $\theta = \theta(\kappa)$ such that for all $x, y \in \hat{\mathcal{X}}_{\mathfrak{U}}$ with $\hat{d}(x, y) \geq \theta$, there exists $U \in \mathfrak{S} - \mathfrak{U}$ such that $d_U(x, y) \geq \kappa$.*

Proof. Let $x, y \in \mathcal{X}$ and let $M = \max_{V \in \mathfrak{S} - \mathfrak{U}} d_V(x, y) + 1$. We may assume $\alpha \geq E$.

We declare $U \in \mathfrak{S}$ to be *relevant* if $d_U(x, y) \geq 100\alpha$. Let \mathfrak{R}^{max} be the set of relevant $T \in \mathfrak{U}$ not properly nested into any relevant element of \mathfrak{U} .

Counting and ordering relevant elements: By Lemma 1.9, there exists $N = N(100\alpha + \kappa)$ so that if $V_1, \dots, V_{N+1} \in \mathfrak{R}^{max}$, then there exists $T \in \mathfrak{S}$ so that $d_T(x, y) \geq 100\alpha + \kappa$ and $V_i \sqsubset T$ for some i . The latter property would ensure that $T \in \mathfrak{S} - \mathfrak{U}$, since \mathfrak{R}^{max} consists of

maximal relevant elements of \mathfrak{U} . Now, if there is such a T , then we are done: we have found $T \in \mathfrak{S} - \mathfrak{U}$ with $d_T(x, y) \geq \kappa$. Hence we may assume that $|\mathfrak{R}^{max}| \leq N$, where $N \geq 1$ depends only on $(\mathcal{X}, \mathfrak{S})$, the constant α , and the desired κ .

By definition, if $U, V \in \mathfrak{R}^{max}$, then $U \triangleleft V$ or $U \perp V$. Hence, let $U \leq V$ if either $U = V$ or $U \triangleleft V$ and $d_U(y, \rho_V^U) \leq E$. As discussed in Section 1.2.2, \leq is a partial ordering on \mathfrak{R}^{max} , and U, V are \leq -incomparable if and only if they are orthogonal. Let V_1, \dots, V_k , with $k \leq N$, be the elements of \mathfrak{R}^{max} , numbered so that $i \leq j$ if $V_i \leq V_j$.

A sequence to estimate $\hat{d}(x, y)$: The idea is to jump from x to P_{V_1} , then from P_{V_1} to P_{V_2} and so on until we get to y . The most “efficient” way of jumping between product regions is described in Lemma 2.6, which justifies the definition of the following sequence of points.¹ Let $x_0 = x, x'_0 = \mathfrak{g}_{P_{V_1}}(x), x_k = \mathfrak{g}_{P_{V_k}}(x), x'_k = \mathfrak{g}_{P_{V_k}}(y)$ and, for $1 \leq i \leq k-1$, let $x_i \in \mathfrak{g}_{P_{V_i}}(P_{V_{i+1}})$ and $x'_i = \mathfrak{g}_{P_{V_{i+1}}}(x_i)$.

There exists $e \in \mathbf{E}_{V_{i+1}}$ and $f \in \mathbf{F}_{V_{i+1}}$ so that $x'_i \in \mathbf{F}_{V_{i+1}} \times \{e\}$ and $x_{i+1} \in \{f\} \times \mathbf{E}_{V_{i+1}}$. Let $z_{i+1} = \mathfrak{g}_{\{f\} \times \mathbf{E}_{V_{i+1}}}(x'_i)$. Observe that $d'(x'_i, z_{i+1}) = 1$.

Bounding $d_U(x_i, x'_i)$: If $d_U(x_i, x'_i) \geq 10^3 \theta E$ for some $U \in \mathfrak{S}$, then $V_i \triangleleft V_{i+1}$ by Lemma 2.6 and [BHS14, Lemma 2.11]. We now bound $d_U(x_i, x'_i)$ for each of the possible types of $U \in \mathfrak{S}$.

First suppose that $U \in \mathfrak{S} - \mathfrak{U}$. Then $U \not\sqsubseteq V_i$ and $U \not\sqsubseteq V_{i+1}$ since \mathfrak{U} is closed under nesting. If $V_i \sqsubseteq U$, then since V_i is $100\alpha > E$ -relevant for x, y , consistency and bounded geodesic image imply that $\rho_U^{V_i}$ lies E -close to any geodesic in $\mathcal{C}U$ from $\pi_U(x)$ to $\pi_U(y)$. The same is true for $\rho_U^{V_{i+1}}$ if $V_{i+1} \sqsubseteq U$. If $V_i \triangleleft U$, then consistency implies that $\rho_U^{V_i}$ lies E -close to $\pi_U(x)$ or $\pi_U(y)$, so that $\rho_U^{V_i}$ again lies E -close to any geodesic from $\pi_U(x)$ to $\pi_U(y)$. Hence, if neither of V_i, V_{i+1} is orthogonal to U , then $d_U(x_i, x'_i) \leq 2(E + \alpha) + M \leq 10^3 \alpha EM$. If $U \perp V_i$, then Lemma 2.6 implies that $d_U(x'_i, x_i) \leq 10^3 \alpha E$. We conclude that $d_U(x'_i, x_i) \leq 10^3 \alpha EM$ whenever $U \in \mathfrak{S} - \mathfrak{U}$.

Next, suppose $U \in \mathfrak{U}$ and $1 \leq i \leq k-1$. If $d_U(x_i, x'_i) > 10^3 \alpha E$, then Lemma 2.6 implies that U is not friendly to V_i or V_{i+1} . Moreover, since $d_U(\rho_U^{V_i}, \rho_U^{V_{i+1}}) \geq 500\alpha E$ by the same lemma, Lemma 2.7 implies that U is relevant, so $U \sqsubseteq U'$ for some \sqsubseteq -maximal relevant $U' \in \mathfrak{U}$. Now, $U \triangleleft V_i, V_{i+1}$ and $U \sqsubseteq U'$, so $U' \notin \{V_i, V_{i+1}\}$. Similarly, we cannot have $U' \perp V_i, V_{i+1}$. Finally, V_i, V_{i+1}, U' are pairwise \sqsubseteq -incomparable, so all are in \mathfrak{R}^{max} and are pairwise \leq -comparable. Note that we can extend \leq to \mathfrak{R}^{max} and observe that \leq is a partial order on $\{U, V_i, V_{i+1}\}$.

If $U' \leq V_i$, then by definition $d_{U'}(y, \rho_{U'}^{V_i}) \leq E$. Since U' is relevant, we have $d_{U'}(x, \rho_{U'}^{V_i}) > E$, so consistency implies that $d_{V_i}(x, \rho_{V_i}^{U'}) \leq E$. Hence $d_{V_i}(\rho_{V_i}^{U'}, y) \geq 50E$. Definition 1.4.(4) implies that $d_{V_i}(\rho_{V_i}^{U'}, \rho_{V_i}^U) \leq E$, so consistency implies that $U < V_i$. Since $U < V_{i+1}$, by transitivity of $<$, we have that $\rho_U^{V_{i+1}}$ coarsely coincides with $\pi_U(y)$. But then $d_U(\rho_U^{V_i}, \rho_U^{V_{i+1}}) \leq 2E$, a contradiction. A similar argument rules out $V_{i+1} < U'$, whence $V_i < U' < V_{i+1}$. However, this contradicts the way we numbered the elements of \mathfrak{R}^{max} . Thus $d_U(x_i, x'_i) \leq 10^3 \alpha E$, as desired.

It remains to bound $d_U(x_0, x'_0)$ for $U \in \mathfrak{U}$ (the case $i = k$ is identical to the case $i = 0$). Suppose that $d_U(x_0, x'_0) \geq 10^3 \alpha E$. The definition of the gate ensures that we cannot have $U \sqsubseteq V_1$ or $U \perp V_1$, so U is not friendly to V_1 . Moreover, $d_U(x_0, \rho_U^{V_1}) \geq 500\alpha E$ since $\rho_U^{V_1}$ coarsely coincides with $\pi_U(x'_0)$. If U is irrelevant, then Lemma 2.7 implies that $d_U(x_0, \rho_U^{V_1}) \leq 200\alpha E$, a contradiction, so U is relevant and $U < V_1$. Also, U is nested in some $U' \in \mathfrak{R}^{max}$. Since $U \sqsubseteq U'$ and U is not friendly to V_1 , we have that $U, U' \triangleleft V_1$ and U', V_1 are $<$ -comparable. Thus $V_1 < U$. Since $\rho_{V_1}^U, \rho_{V_1}^{U'}$ coarsely coincide, we have that $\rho_{V_1}^{U'}$ is far from x_0 , so that $U' < V_1$, which is impossible. Hence $d_U(x_0, x'_0) \leq 10^3 \alpha E$.

¹It seems natural to take $x_i = \pi_{V_i}(y), x'_i = \pi_{V_{i+1}}(x)$, but in certain situations this would create extraneous distance formula terms between x_i, x'_i , namely when there exists $U \sqsubseteq V_i, V_{i+1}$.

For any $\mu' \geq 10^3 \alpha E$, we have shown that $d_U(x_i, x'_i) \leq \mu' M$ when $U \in \mathfrak{S} - \mathfrak{U}$ and $d_U(x_i, x'_i) \leq \mu'$ when $U \in \mathfrak{U}$, for $0 \leq i \leq k$.

Bounding $\hat{d}(x, y)$: We have produced a uniform constant μ' so that $d_U(x_i, x'_i) \leq \mu' M$ when $U \in \mathfrak{S} - \mathfrak{U}$ and $d_U(x_i, x'_i) \leq \mu'$ when $U \in \mathfrak{U}$, for $0 \leq i \leq k$. Hence, by the distance formula (Theorem 1.13) with threshold $\mu' + 1$, we have $d'(x_i, x'_i) \leq \mu N M$ for some uniform μ . Thus $\sum_{i=0}^k d'(x_i, x'_i) \leq (k+1)\mu N M \leq 2\mu N^2 M$. (Recall that N depends only on $(\mathcal{X}, \mathfrak{S})$, the set \mathfrak{U} , and the input κ .)

Now,

$$\begin{aligned} \hat{d}(x, y) &\leq \sum_{i=0}^k d'(x_i, x'_i) + \sum_{i=0}^{k-1} [d'(x'_i, z_{i+1}) + \hat{d}(z_{i+1}, x_{i+1})] \\ &\leq 2\mu N^2 M + N + \sum_{i=0}^{k-1} \hat{d}(z_{i+1}, x_{i+1}). \end{aligned}$$

Fix $0 \leq i \leq k-1$, let $\mathbf{E} = \mathbf{E}_{V_{i+1}}$ and $\mathbf{F} = \mathbf{F}_{V_{i+1}}$, for convenience. Consider the hierarchical space $(\mathbf{E}, \mathfrak{S}_{V_{i+1}}^\perp)$, where $\mathfrak{S}_{V_{i+1}}^\perp$ consists of all those $U \in \mathfrak{S}$ with $U \perp V_{i+1}$ together with some $A \in \mathfrak{S}$ such that $A \sqsubset S$ and each U orthogonal to V_{i+1} satisfies $U \sqsubseteq A$. Let $\mathfrak{U}_{\mathbf{E}} = \mathfrak{U} \cap \mathfrak{S}_{V_{i+1}}^\perp$ and consider the factored space $(\hat{\mathbf{E}}_{\mathfrak{U}_{\mathbf{E}}}, \mathfrak{S}_{V_{i+1}}^\perp - \mathfrak{U}_{\mathbf{E}})$, whose metric we denote $\hat{d}_{\mathbf{E}}$. Observe that $\hat{d}(z_{i+1}, x_{i+1}) \leq \epsilon \hat{d}_{\mathbf{E}}(z_{i+1}, x_{i+1}) + \epsilon$ for some uniform ϵ , since $\mathbf{E} \rightarrow \mathcal{X}$ is a uniform quasi-isometric embedding and any two points lying on a parallel copy of some \mathbf{F}_U in \mathbf{E} also lie on such a parallel copy in \mathcal{X} .

Now, by induction on complexity, there exists a “uniqueness function” $f : \mathbb{N} \rightarrow \mathbb{N}$, independent of i , so that $(\hat{\mathbf{E}}_{\mathfrak{U}_{\mathbf{E}}}, \mathfrak{S}_{V_{i+1}}^\perp - \mathfrak{U}_{\mathbf{E}})$ has the following property: if $e, e' \in \mathbf{E}$, then

$$\hat{d}(e, e') \leq \epsilon f \left(\max_{U \in \mathfrak{S}_{V_{i+1}}^\perp - \mathfrak{U}_{\mathbf{E}}} d_U(e, e') \right) + \epsilon.$$

Indeed, in the base case, either $\mathfrak{S}_{V_{i+1}}^\perp - \mathfrak{U}_{\mathbf{E}} = \emptyset$, and \mathbf{E} is uniformly bounded in \mathcal{X} (and hence its \hat{d} -diameter is uniformly bounded) or $\mathfrak{U}_{\mathbf{E}} = \emptyset$ and $\hat{d}_{\mathbf{E}}$ coarsely coincides with d on \mathbf{E} , whence f exists by uniqueness in \mathbf{E} (with metric d and HS structure $\mathfrak{S}_{V_{i+1}}^\perp$).

Claim 1. There exists $\eta = \eta((\mathcal{X}, \mathfrak{S}))$ such that for all $U \in \mathfrak{S} - \mathfrak{U}$ and $1 \leq i \leq k-1$, there exists $x_i \in \mathfrak{g}_{P_{V_i}}(P_{V_{i+1}})$ so that $\pi_U(x_i)$ lies η -close to a geodesic from $\pi_U(x)$ to $\pi_U(y)$.

Proof of Claim 1. Theorem 1.14 provides a D -discrete D -hierarchy path γ joining y to x , where D depends only on $(\mathcal{X}, \mathfrak{S})$. We may assume that $D \leq \alpha$, since α was chosen in advance in terms of $(\mathcal{X}, \mathfrak{S})$ only. The proof of Proposition 5.16 of [BHS15] (which does not use hyperbolicity of the various $\mathcal{C}U$), provides a constant η' so that, for each i , there exists a maximal subpath γ_i of γ lying in $\mathcal{N}_{\eta'}(\mathbf{P}_{V_i})$, with initial point x'_i . Moreover, $\pi_U(x'_i)$ uniformly coarsely coincides with $\rho_U^{V_i}$ when $V_i \sqsubset U$ or $V_i \pitchfork U$. Let $y_i \in \mathbf{P}_{V_{i+1}}$ lie η'' -close to the terminal point of γ_{i+1} . We claim that $d(\mathfrak{g}_{\mathbf{P}_{V_i}}(y_i), x'_i)$ is uniformly bounded. Indeed, by definition $\pi_U(\mathfrak{g}_{\mathbf{P}_{V_i}}(y_i))$ coarsely coincides with $\rho_U^{V_i}$, and hence with $\pi_U(x)$, when $U \pitchfork V_i$ or $V_i \sqsubseteq U$, and coincides with $\pi_U(y_i)$ when $U \sqsubseteq V_i$ or $U \perp V_i$. Our choice of x'_i, y_i ensures that $d_U(x'_i, y_i)$ is uniformly bounded for such U , so our claim follows from the distance formula. Taking $x_i = \mathfrak{g}_{\mathbf{P}_{V_i}}(y_i)$ completes the proof, since x'_i , and hence x_i , lies uniformly close to any geodesic from $\pi_U(x)$ to $\pi_U(y)$ in any δ -hyperbolic $\mathcal{C}U$, by the definition of a hierarchy path. \blacksquare

We now choose specific values of x_i, x'_i, z_i satisfying the above defining conditions. First, as before, $x_0 = x, x'_0 = \mathfrak{g}_{\mathbf{P}_{V_1}}(x)$, while $x_k = \mathfrak{g}_{\mathbf{P}_{V_k}}(y)$ and $x'_k = y$. For $1 \leq i \leq k-1$, let x_i

be a point provided by Claim 1. Then let $x'_i = \mathfrak{g}_{P_{V_{i+1}}}(x_i)$ for $1 \leq i \leq k-1$, as before, and define the points z_i as above.

Claim 2. There exists a function $f': \mathbb{N} \rightarrow \mathbb{N}$, independent of i , so that $\mathbf{d}_U(z_{i+1}, x_{i+1}) \leq f'(M)$ for all $U \in \mathfrak{S}_{V_{i+1}}^\perp - \mathfrak{U}_{\mathbf{E}}$.

Proof of Claim 2. Let $U \in \mathfrak{S}_{V_{i+1}}^\perp - \mathfrak{U}_{\mathbf{E}}$. By the definition of gates, for $1 \leq i \leq k-1$, we have $\mathbf{d}_U(z_{i+1}, x_i) \leq \alpha$ so, since $\pi_U(x_i)$ lies η -close to a geodesic from $\pi_U(x)$ to $\pi_U(y)$, we have $\mathbf{d}_U(z_{i+1}, \{x, y\}) \leq \eta + \alpha$. Likewise, $\mathbf{d}_U(x_{i+1}, \{x, y\}) \leq \eta$. Hence $\mathbf{d}_U(z_{i+1}, x_{i+1}) \leq M + 2\eta + \alpha$. \blacksquare

Claim 2 and the above discussion imply that, if $M \leq \kappa$, we have

$$\hat{\mathbf{d}}(x, y) \leq 2\mu N^2 \kappa + N + \epsilon N f(f'(\kappa)) + N\epsilon,$$

which completes the proof. \square

2.1. $\mathcal{C}(\ast)$ as a coarse intersection graph. We conclude this section by highlighting a particularly interesting application of Proposition 2.4, one of the tools we developed for proving the results about asymptotic dimension. ²

Corollary 2.9. *Given a relatively HHS $(\mathcal{X}, \mathfrak{S})$, the space $\hat{\mathcal{X}}_{\mathfrak{S}-\{S\}}$ is quasi-isometric to $\pi_S(\mathcal{X}) \subseteq \mathcal{CS}$, where $S \in \mathfrak{S}$ is \sqsubseteq -maximal.*

Proof. By Proposition 2.4, $(\hat{\mathcal{X}}_{\mathfrak{S}-\{S\}}, \{S\})$ is a hierarchically hyperbolic space, and the claim follows from the distance formula (Theorem 1.13). \square

Remark 2.10 (Coning off \mathbf{P}_U). If we had constructed $\hat{\mathcal{X}}_{\mathfrak{S}-\{S\}}$ by “coning off” \mathbf{P}_U for each $U \in \mathfrak{S} - \{S\}$, instead of coning off each parallel copy of each \mathbf{F}_U , then Corollary 2.9 would continue to hold.

In many examples of interest, π_S is coarsely surjective, so that Corollary 2.9 yields a quasi-isometry $\hat{\mathcal{X}}_{\mathfrak{S}-\{S\}} \rightarrow \mathcal{CS}$. Moreover, if $(\mathcal{X}, \mathfrak{S})$ is an HHS, then (as described in [BHS15, DHS16]), \mathcal{X} admits an HHS structure obtained by replacing each \mathcal{CU} with a hyperbolic space quasi-isometric to $\pi_U(\mathcal{X})$, so in particular \mathcal{CS} becomes quasi-isometric to the space obtained by coning off each parallel copy of each $\mathbf{F}_U, U \neq S$. If, as is the case for hierarchically hyperbolic groups $(\mathcal{G}, \mathfrak{S})$, the parallel copies of the various \mathbf{F}_U coarsely cover \mathcal{X} , this provides a hierarchically hyperbolic structure in which \mathcal{CS} is a coarse intersection graph of the set of \mathbf{F}_U with $U \sqsubset S$ for which there is no V with $U \sqsubset V \sqsubset S$. This is a coarse version of what happens, for example, when \mathcal{X} is a CAT(0) cube complex with a factor system and we can take \mathcal{CS} to be the *contact graph* of \mathcal{X} (see [Hag14, BHS14]).

3. ASYMPTOTIC DIMENSION OF THE \mathcal{CU}

In this section $(\mathcal{X}, \mathfrak{S})$ is a relatively hierarchically hyperbolic space with the additional property that \mathcal{X} is a *uniformly locally-finite discrete geodesic space*, i.e.,

- (1) there exists $r_0 > 0$ so that $\mathbf{d}(x, y) \geq r_0$ for all distinct $x, y \in \mathcal{X}$;
- (2) there is a function $p: [0, \infty) \rightarrow [0, \infty)$ so that $|B(x, r)| \leq p(r)$ for all $x \in \mathcal{X}$;
- (3) there exists r_1 so that for all $x, y \in \mathcal{X}$, there exists n and $\gamma: \{0, n\} \rightarrow \mathcal{X}$ so that $\gamma(0) = x, \gamma(n) = y$, and $\mathbf{d}(x, y) = \sum_{i=0}^{n-1} \mathbf{d}(\gamma(i), \gamma(i+1))$, and $\mathbf{d}(\gamma(i), \gamma(i+1)) \leq r_1$ for all i .

If \mathcal{X} satisfies (1) and (3) (but not necessarily (2)), then \mathcal{X} is a (r_0, r_1) -*discrete geodesic space*.

The following notion is motivated by work of Bowditch; see [Bow08, Section 3].

²A proof of the results in this subsection appeared in the first version of [BHS15] using techniques which we have now generalized to prove Proposition 2.4.

Definition 3.1 (Tight space). The δ -hyperbolic space F is (C, K) -tight if there exists a map $\beta: F^2 \rightarrow 2^F$ so that:

- (1) for every $x, y \in F$, we have $\mathbf{d}_{Haus}([x, y], \beta(x, y)) \leq C$, where $[x, y]$ is any geodesic from x to y ;
- (2) for every $x, y, z \in F$ with y on a geodesic from x to z , if $\mathbf{d}(y, \{x, z\}) \geq r + C$ for some $r \in \mathbb{R}^+$ then $B(y, 2\delta + 2C) \cap \bigcup_{x' \in B(x, r), z' \in B(z, r)} \beta(x', z')$ has cardinality at most K .

Notice that for x, y, z, r as in 2, if $x' \in B(x, r), z' \in B(z, r)$ then $\beta(x', z')$ intersects $B(y, 2\delta + 2C)$ by δ -hyperbolicity and part 1 of Definition 3.1.

Let ξ denote the complexity of \mathfrak{S} , i.e., the length of a longest \sqsubseteq -chain. By Definition 1.4.(5), $\xi < \infty$. The aim of this section is to prove:

Theorem 3.2 (Existence of tight geodesics). *Let $(\mathcal{X}, \mathfrak{S})$ be a δ -relatively hierarchically hyperbolic space, where \mathcal{X} is a uniformly locally-finite discrete geodesic space. Suppose moreover that $\pi_U: \mathcal{X} \rightarrow \mathcal{CU}$ is uniformly coarsely surjective, where U varies over all elements of \mathfrak{S} with \mathcal{CU} a δ -hyperbolic space. Then there exist $C, K \geq 0$ so that \mathcal{CU} is (C, K) -tight for every $U \in \mathfrak{S}$ for which \mathcal{CU} is δ -hyperbolic.*

Proof. Throughout this proof we use the identification of \mathcal{CS} with the coned-off space $\hat{\mathcal{X}}$, as established in Proposition 2.4. Our assumption on coarse surjectivity of the projections π_U implies that, for each $U \in \mathfrak{S}$ with \mathcal{CU} a δ -hyperbolic space, we may (by an initial change in the constants from Definition 1.4), assume that π_U is actually surjective.

Fix a constant D , as provided by Theorem 1.14, so that every pair of points of \mathcal{X} can be joined by a D -hierarchy path. For S the \sqsubseteq -maximal element of \mathfrak{S} , we will show that \mathcal{CS} is (C, K) -tight, where C, K depend only on D, E , the complexity ξ of \mathfrak{S} , and the function p which quantifies the local-finiteness. To see that this suffices, recall that for each $U \in \mathfrak{S}$, there is a hierarchically quasiconvex subspace \mathbf{F}_U with a relatively hierarchically hyperbolic structure $(\mathbf{F}_U, \mathfrak{S}_U)$ in which U is \sqsubseteq -maximal.

For $M \geq 0$ and $x, y \in \mathcal{X}$, define

$$\beta_M(x, y) = \{z \in \mathcal{X} : \hat{\mathbf{d}}(z, [x, y]) \leq M, \mathbf{d}_U(z, \{x, y\}) \leq M \forall U \in \mathfrak{S} - \{S\}\}.$$

For sufficiently large M , the map $\beta_M: (\mathcal{CS})^2 \rightarrow 2^{\mathcal{CS}}$ satisfies property (1), by the definition and Claim 1 below.

Claim 1. For each sufficiently large M there exists K_1 so that for each $x, y \in \mathcal{X}$ and $z \in [x, y]$ we have $\hat{\mathbf{d}}(z, \beta_M(x, y)) \leq K_1$.

Proof. We distinguish two cases. First, suppose that there exists $U \in \mathbf{Rel}(x, y, 10E)$ for which $\hat{\mathbf{d}}(\rho_S^U, z) \leq 10DE$, and take a \sqsubseteq -maximal U with such property. Then consider $z' = \mathbf{g}_{P_U}(x)$. Clearly, $\hat{\mathbf{d}}(z, z')$ is uniformly bounded, and we now show $z' \in \beta_M(x, y)$, provided that M is large enough. In order to do so, we uniformly bound $\mathbf{d}_V(\{x, y\}, z')$ for each $V \in \mathfrak{S} - \{S\}$. If V is either nested into U or orthogonal to U , then we are done by the definition of gate. Otherwise, $\pi_V(z')$ coarsely coincides with ρ_V^U , so we have to show that either $\pi_V(x)$ or $\pi_V(y)$ coarsely coincides with ρ_V^U . Suppose by contradiction that this is not the case. If $U \pitchfork V$, then consistency implies that $\pi_U(x), \pi_U(y)$ both coarsely coincide with ρ_V^U , so that $\mathbf{d}_U(x, y) \leq 3E$ contradicting the choice of U . If $U \not\sqsubseteq V$, then any geodesic from $\pi_V(x)$ to $\pi_V(y)$ stays far from ρ_V^U since, by maximality of U , we have $\mathbf{d}_V(x, y) < 10E$. In particular, by bounded geodesic image and consistency we have $\mathbf{d}_U(x, y) < 10E$, a contradiction.

Suppose now that there does not exist any $U \in \mathbf{Rel}(x, y, 10E)$ so that $\hat{\mathbf{d}}(\rho_S^U, z) \leq 10DE$. Consider a hierarchy path γ from x to y . Since π_S is coarsely Lipschitz, there exists $z' \in \gamma$ so that $\hat{\mathbf{d}}(z, z') \leq 5DE$. We claim $z' \in \beta_M(x, y)$ for sufficiently large M . In fact, for any $U \in \mathfrak{S} - \{S\}$ (notice that if such U exists then $\hat{\mathcal{X}}$ is δ -hyperbolic) either $\hat{\mathbf{d}}(\rho_S^U, z') \leq 5E$, so




FIGURE 1. The second case in the proof of Claim 1 when each relevant U has ρ_S^U far from z . The subcases in the proof correspond to the $U_1, U_2 \in \mathfrak{S}$ shown.

$\hat{d}(\rho_S^U, z) \leq 5DE + 5E \leq 10DE$, in which case U is irrelevant and we are done by hypothesis, or ρ_S^U lies $5E$ -far from a geodesic from $\pi_S(z')$ to one of $\pi_S(x), \pi_S(y)$. In this case, we can apply bounded geodesic image to conclude. See Figure 1. \blacksquare

We now prove Property 2. Let $x, y, z \in \mathcal{X}$ with y on a (discrete) \hat{d} -geodesic from x to z , and suppose $\hat{d}(y, \{x, z\}) \geq r + K_1$ for some $r \geq K_1 + 2\delta + 100E$. Let

$$T = B^{\hat{\mathcal{X}}}(y, 2\delta + 2K_1) \cap \left(\bigcup_{x' \in B^{\hat{\mathcal{X}}}(x, r), z' \in B^{\hat{\mathcal{X}}}(z, r)} \beta_M(x', z') \right).$$

Moreover, let $\mathcal{U} = \mathcal{U}(x, y, z)$ be the set of all $U \in \mathfrak{S}$ satisfying the following conditions:

- $d_U(x, z) \geq 10M + 10^3E$,
- $U \neq S$ and $d_S(\rho_S^U, y) \leq 2\delta + 2K_1 + 50E$,
- whenever $V \in \mathfrak{S}$ satisfies $d_V(x, z) \geq 10M + 10^3E$ and $U \subsetneq V \subsetneq S$, we have $d_V(\rho_V^U, \{x, z\}) \leq 100E + 10M$.

We will see in what follows that for $y' \in T$ it is sufficient to have information about $\pi_U(y')$ for $U \in \mathcal{U}$ to coarsely reconstruct all $\pi_V(y')$. Moreover, we will bound the cardinality of \mathcal{U} .

For $y' \in T$, denote $\mathcal{U}(y') = \{U \in \mathcal{U} : d_U(x, y') \leq M + 10E\}$, and let \mathcal{B} be the set of all subsets of \mathfrak{S} of the form $\mathcal{U}(y'), y' \in T$.

Claim 2. There is $K_3 = K_3(E, K_1, M)$ with $|\{y' \in T : \mathcal{U}(y') = \mathcal{U}_i\}| \leq p(K_3)$ for all $\mathcal{U}_i \in \mathcal{B}$.

Proof of Claim 2. Fix $y', y'' \in T$ with $\mathcal{U}(y') = \mathcal{U}(y'')$. We bound $d(y', y'')$ by bounding $d_U(y', y'')$ for all $U \in \mathfrak{S}$.

Let $U \in \mathfrak{S}$. If $U \in \mathcal{U}$, then $d_U(y', y'') \leq 2M + 10E$ since $U(y') = U(y'')$ (and the fact that each of $\pi_U(y'), \pi_U(y'')$ coarsely coincides with either $\pi_U(x)$ or $\pi_U(z)$ in $\mathcal{C}U$). We now analyze the other cases, but for technical reasons we change the constants from the definition of \mathcal{U} .

If $U = S$, then $d_U(y', y'') \leq 4\delta + 4K_1$. Assume $U \neq S$ from now on.

If $d_U(x, z) < 10M + 150E$, then $d_U(y', y'') \leq 20M + 200E$.

If $d_S(\rho_S^U, y) > 2\delta + 2K_1 + 10E$, then no geodesic from y' to y'' in $\hat{\mathcal{X}}$ passes E -close to ρ_S^U , so consistency and bounded geodesic image again yield $d_U(y', y'') \leq 10E$.

Finally, suppose that $d_U(x, z) \geq 10M + 150E$ and $d_S(\rho_S^U, y) \leq 2\delta + 2K_1 + 10E$, but that there exists $V \in \mathfrak{S}$ satisfying $d_V(x, z) \geq 10M + 10^3E$, $U \subsetneq V \subsetneq S$ and $d_V(\rho_V^U, \{x, z\}) > 50E + 10M$. Consider a \sqsubseteq -maximal V with this property. We claim that $V \in \mathcal{U}$. In fact, ρ_S^V coarsely coincides with ρ_S^U , yielding the second condition in the definition of \mathcal{U} . Moreover, for any $W \in \mathfrak{S}$ with $V \subsetneq W \subsetneq S$, we have $d_W(\rho_W^V, \{x, z\}) \leq 100E + 10M$, for otherwise V would not be maximal (we are once again using that ρ_W^V coarsely coincides with ρ_W^U).

Now, since $\mathcal{U}(y') = \mathcal{U}(y'')$ we have that $\pi_V(y'), \pi_V(y'')$ are both close to one of $\pi_V(x), \pi_V(z)$. Hence, $\pi_U(y')$ must $10E$ -coarsely coincide with $\pi_U(y'')$ because geodesics in $\mathcal{C}V$ from $\pi_V(y')$ to $\pi_V(y'')$ stay E -far from ρ_V^U .

We conclude that for all $U \in \mathfrak{S}$, we have, say, $d_U(y', y'') \leq 500ME\delta K_1$, so the distance formula (Theorem 1.13) provides a uniform $K_3 = K_3(M, E, K_1)$ so that $d_{\mathcal{X}}(y', y'') \leq K_3$, and the claim follows from the definition of p . \blacksquare

Claim 3. There exists $K_4 = K_4(E, M)$ so that $|\mathcal{B}| \leq 2^{K_4^\xi}$.

Proof. By definition, if $U \in \mathcal{U}$, then $d_S(\rho_S^U, y) \leq 2\delta + 2K_1 + 50E$. Choose $a, b \in \mathcal{X}$ such that a is E -close to a $\hat{\mathcal{X}}$ -geodesic from y to z and satisfies $\hat{d}(y, a) \in [2\delta + 2K_1 + 60E, 2\delta + 2K_1 + 80E]$, while b is E -close to a $\hat{\mathcal{X}}$ -geodesic from x to y and satisfies $\hat{d}(y, b) \in [2\delta + 2K_1 + 60E, 2\delta + 2K_1 + 80E]$. Then ρ_S^U does not lie E -close to a geodesic from a to z or from x to b , so consistency and bounded geodesic image yield $d_U(a, z) \leq E$ and $d_U(x, b) \leq E$. Thus $d_U(b, a) \geq 10M + 98E \geq E$, so Lemma 1.9 yields $K_5 = K_5(M, E)$ so that there are at most K_5 such U that are \sqsubseteq -maximal.

Fix $U \in \mathcal{U} - \{S\}$ to be \sqsubseteq -maximal and consider all $V \in \mathfrak{S}$ such that $V \subsetneq U$ and such that the following conditions are satisfied:

- $d_V(x, z) \geq 10M + 10^3E$,
- $d_S(\rho_S^V, y) \leq 2\delta + 2K_1 + 50E$,
- $d_U(\rho_U^V, x) \leq 100E + 10M$.

(Note that any $V \in \mathcal{U}$ nested into U satisfies either these conditions or the same conditions with z replacing x in the third one.) By partial realization and the first and third conditions above, there exists $a_U \in \mathcal{X}$ such that $d_U(a_U, x) \leq 10M + 200E$ and ρ_V^U fails to come E -close to every geodesic from $\pi_U(x)$ to $\pi_U(a_U)$ for all V as above. Hence, for each V as above, bounded geodesic image and consistency imply that $d_U(a_U, x) \leq 10M + 200E$ and $d_V(a_U, x) \geq 10M + 9E$. Let \mathcal{V}_0 be the set of all \sqsubseteq -maximal $V \subsetneq U$ contained in \mathcal{U} . By Lemma 1.9 and maximality, there exists $K_6 = K_6(10M + 200E)$ so that $|\mathcal{V}_0| \leq K_6$.

For each $V \in \mathcal{V}_0 \cap \mathcal{U}$, choose a_V as above, so that $d_V(a_V, x) \leq 10M + 200E$ and, any element W of \mathcal{U} properly nested into V and satisfying $d_V(\rho_V^W, x) \leq 100E + 10M$, we have $d_W(a_V, x) \geq 10M + 9E$. Then, exactly as above, we find that there are at most K_6 such W that are \sqsubseteq -maximal and properly nested in V . Proceeding inductively, we see that there are at most $K_6^{\xi-1}$ elements T of \mathcal{U} satisfying $d_U(\rho_U^T, x) \leq 100E + 10M$ and properly nested into U , while an identical discussion bounds the set of $T \in \mathcal{U}$ properly nested in U and satisfying $d_U(\rho_U^T, z) \leq 100E + 10M$. Hence $|\mathcal{U}| \leq 2K_5K_6^{\xi-1}$, so, letting $K_4 = \max\{2K_5, K_6\}$, we get $|\mathcal{B}| \leq 2^{K_4^\xi}$. \blacksquare

Claim 2 and Claim 3 together imply that $|T| \leq p(K_3) \cdot 2^{K_4^\xi}$ for uniform K_4 , so $\mathcal{C}S$ is $(K_1, p(K_3) \cdot 2^{K_4^\xi})$ -tight. \square

Corollary 3.3. *Let \mathcal{X} be a uniformly locally-finite discrete geodesic space and let $(\mathcal{X}, \mathfrak{S})$ be δ -relatively hierarchically hyperbolic. Suppose moreover that $\pi_U : \mathcal{X} \rightarrow \mathcal{C}U$ is uniformly coarsely surjective, where U varies over all elements of \mathfrak{S} with $\mathcal{C}U$ a δ -hyperbolic space. Then there exist $\lambda, \mu = \lambda(\delta, D, E, p), \mu(\delta, D, E, p)$, independent of ξ , so that $\text{asdim } \mathcal{C}U \leq \lambda \cdot 2^{\mu\xi}$ uniformly, whenever $U \in \mathfrak{S}$ has the property that $\mathcal{C}U$ is δ -hyperbolic.*

Proof. By Theorem 3.2, there exist C, λ', μ so that $\mathcal{C}U$ is $(C, \lambda' \cdot 2^{\mu\xi})$ -tight for each $U \in \mathfrak{S}$ with $\mathcal{C}U$ a δ -hyperbolic space. We now argue roughly as in the proof of [BF07, Theorem 1].

Fix $r \in \mathbb{N}, x_0 \in \mathcal{C}U$, and $\ell \in \mathbb{N}$ with $\ell \geq 10\delta + C$. For each $n \geq 1$, let

$$A_n = \{x \in \mathcal{C}U : 10(n-1)(r+\ell) \leq d_U(x, x_0) \leq 10n(r+\ell)\},$$

so that $\cup_n A_n = \mathcal{CU}$. Let $S_n = \{x \in \mathcal{CU} : d_U(x, x_0) = 10n(r + \ell)\}$. Given $n \geq 3$, define subsets $B_n^i \subset A_n$ as follows: and for each $s_i \in S_{n-2}$ and $x \in A_n$, set $x \in B_n^i$ if and only if there exists $\beta(x, x_0)$ lying at Hausdorff distance C from a geodesic joining x_0 to x that passes through s_i . For $n \in \{1, 2\}$, define $B_n^1 = A_n$. Then the sets $\{B_n^i\}$ cover \mathcal{CU} . Observe that $\text{diam } B_n^i \leq 100C(r + \ell)$ for each n, i .

Hence $\{B_n^i\}$ is a uniformly bounded cover of \mathcal{CU} . We now check that its $\frac{r}{2}$ -multiplicity is at most $2\lambda' \cdot 2^{\mu^\varepsilon}$, from which the claim follows. Our claim is a consequence of the following statement: for $n \geq 1$, there are $\leq \lambda' \cdot 2^{\mu^\varepsilon}$ elements of $\{B_n^i\}$ intersecting any $\frac{r}{2}$ -ball \mathcal{B} in \mathcal{CU} .

Indeed, the claim is clear when $n \leq 2$, so assume $n \geq 3$. Choose distinct B_n^i, B_n^j intersecting \mathcal{B} , with the intersections respectively containing points y_i, y_j . Hence we have geodesics $[x_0, y_i], [x_0, y_j]$, respectively joining x_0 to y_i, y_j and passing through s_i, s_j , and paths $\beta_i = \beta(x_0, y_i)$ and $\beta_j = \beta(x_0, y_j)$ at Hausdorff distance $\leq C$ from $[x_0, y_i], [x_0, y_j]$. Choose $s'_i \in \beta_i, s'_j \in \beta_j$ with $d_U(s_i, s'_i), d_U(s_j, s'_j) \leq C$. Since $d_U(y_i, y_j) \leq r$, because $y_i, y_j \in \mathcal{B}$, we see that s'_i, s'_j both lie in $\bigcup_{x' \in B(x_0, r), z' \in B(y_i, r)} \beta(x', z')$.

Now, s_i lies on a geodesic from x_0 to y_i and $d_U(s_i, \{x_0, y_i\}) \geq r + C$ by definition. Hyperbolicity and the definition of S_{n-2} imply that $d_U(s_i, s_j) \leq 2\delta$, whence $d_U(s'_j, s_i), d_U(s'_i, s_i) \leq 2\delta + 2C$, i.e., $s'_i, s'_j \in B(s_i, \delta + 2C) \cap \bigcup_{x' \in B(x_0, r), z' \in B(y_i, r)} \beta(x', z')$. This intersection is bounded by Property (2) of tightness, thus completing the proof. \square

4. ASYMPTOTIC DIMENSION OF BALL-PREIMAGES UNDER $\psi: \mathcal{X} \rightarrow \hat{\mathcal{X}}$

Fix a relatively hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$. Let \mathcal{O} be the set of totally orthogonal subsets of $\mathfrak{U}_1 \subseteq \mathfrak{S}$. We set $\text{asdim } \mathcal{O}$ to be the minimal uniform asymptotic dimension of the collection $\{\Pi_{V \in \mathfrak{O}} \mathbf{F}_V : \mathfrak{O} \in \mathcal{O}\}$. Our goal is to prove:

Proposition 4.1. *For every $R \geq 0$, there exists a hierarchical quasiconvexity function k and a control function f , so that for every $x_0 \in \hat{\mathcal{X}}$ there exists $A \subset \mathcal{X}$ satisfying:*

- A is k -hierarchically quasiconvex in \mathcal{X} ,
- $A \supseteq \psi^{-1}(B(x_0, R))$,
- $\text{asdim } A \leq \text{asdim } \mathcal{O}$ with control function f .

Proof. For convenience, fix $x_0 \in \hat{\mathcal{X}}$ and stipulate that all constants, hierarchical quasiconvexity functions, and control functions below are chosen independently of x_0 . Fix D large enough that any pair of points in \mathcal{X} are joined by a (D, D) -quasigeodesic.

Set $A_0 = \{x_0\}$. We will inductively construct A_n, k_n, f_n so that

- (1) $A_n \subseteq \mathcal{X}$ is k_n -hierarchically quasiconvex,
- (2) $\text{asdim } A_n \leq \text{asdim } \mathcal{O}$ with control function f_n ,
- (3) for each $n \geq 1$, the set A_n contains $\mathcal{N}_{10D}(A_{n-1})$,
- (4) for each $n \geq 1$ and $U \in \mathfrak{U}_1$, the set A_n contains each parallel copy of \mathbf{F}_U which intersects A_{n-1} .

Assuming for the moment that we can construct such A_n , then the proposition follows. Indeed, given $R \geq 1$, we now show that the set A_n contains $\psi^{-1}(B(x_0, R))$ for $n \geq 10D^2R$. Let $x \in \psi^{-1}(B(x_0, R))$, whence, by definition of d' , there exists a sequence $x_0, \dots, x_k = x$ for which $\sum_{i=1}^{k-1} d'(x_i, x_{i+1}) \leq 2R$. Using that \mathcal{X} is a quasigeodesic space, we can suitably interpolate between consecutive points in the sequence and find another sequence $x_0 = y_0, \dots, y_{k'} = x$ so that $\sum_{i=1}^{k'-1} d'(y_i, y_{i+1}) \leq 10D^2R$ and for each i either y_i, y_{i+1} lie in a common \mathbf{F}_U for $U \in \mathfrak{U}_1$ or $d(y_i, y_{i+1}) \leq 10D$. It is readily shown inductively that y_i lies in A_i for each i , so that in particular $x \in A_n$, as required.

Fix $n \geq 1$ and assume we have constructed A_{n-1} with the desired properties; we now construct A_n . Let A'_{n-1} be the $10D$ -neighborhood of A_{n-1} . Note that, for suitable k'_{n-1}, f'_{n-1}

depending on k_{n-1} , f_{n-1} , and D , the space A'_{n-1} is k'_{n-1} -hierarchically quasiconvex and has asymptotic dimension at most $\text{asdim } \mathcal{O}$ with control function f'_{n-1} .

Let M be a constant to be chosen later. We say $U \in \mathfrak{U}_1$ is *admissible* if whenever $V \in \mathfrak{S}$ satisfies either $U \pitchfork V$ or $U \sqsubset V$ we have $d_V(\rho_V^U, A_{n-1}) \leq M$. The totally orthogonal collection $\mathcal{U} \subseteq \mathfrak{U}_1$ is admissible if each $U \in \mathcal{U}$ is admissible. For an admissible collection \mathcal{U} we define $B_{\mathcal{U}} \subseteq \mathcal{X}$ to be the set of all $x \in \mathcal{X}$ so that

- $d_V(x, A'_{n-1}) \leq M$ for each $V \perp \mathcal{U}$;
- $d_V(x, \rho_V^U) \leq M$ whenever $V \in \mathfrak{S}$ and $U \in \mathcal{U}$ satisfy $U \pitchfork V$ or $U \sqsubset V$.

(Roughly, $B_{\mathcal{U}}$ is the set of partial realization points for \mathcal{U} whose projections lie close to A'_{n-1} except possibly in $\mathcal{C}U$ for $U \in \mathcal{U}$.)

Since A_{n-1} is hierarchically quasiconvex, B_{\emptyset} coarsely coincides with A_{n-1} . Let $A_n = A'_{n-1} \cup \bigcup_{\mathcal{U}} B_{\mathcal{U}}$, where the union is taken over all admissible collections \mathcal{U} . The desired properties of A_n , for suitable k_n, f_n , will be checked in the following claims.

Claim 1. If $U \in \mathfrak{U}_1$ has the property that $\mathbf{F}_U \times \{e\}$ intersects A_{n-1} for some $e \in \mathbf{E}_U$, then $\mathbf{F}_U \times \{e\} \subseteq A_n$.

Proof. For U as in the statement and M large enough, $\text{diam}_V(\rho_V^U \cup \pi_V(\mathbf{F}_U)) \leq M$ whenever $U \pitchfork V$ or $U \sqsubset V$, by the definition of \mathbf{F}_U . Hence, since π_V is coarsely Lipschitz and $\pi_V(\mathbf{F}_U)$ intersects $\pi_V(A_{n-1})$, we also have that $d_V(\rho_V^U, A_{n-1}) \leq M$. Thus U is admissible. Any point x in a parallel copy of \mathbf{F}_U which intersects A_{n-1} is readily seen to be in $B_{\{U\}} \subseteq A_n$, as required. \blacksquare

Claim 2. For each admissible $U \in \mathfrak{U}_1$, the projections $\pi_U(\mathcal{X})$ and $\pi_U(A_n)$ coarsely coincide.

Proof. Suppose that U is admissible and let $p \in \pi_U(\mathcal{X})$. Fix $y \in A_{n-1}$ and define the tuple \vec{b} in the following way. Let $b_V = \pi_V(y)$ if $V \perp U$, $b_V = \rho_V^U$ if $V \in \mathfrak{S}$ satisfies $U \pitchfork V$ or $U \sqsubset V$, and $b_U = p$. It is easy to check that \vec{b} is consistent, using the definitions and Proposition 1.5. Hence, for M large enough, realization (Theorem 1.12) provides us with a point $x \in \mathcal{X}$ which, by the definition of \vec{b} , is contained in $B_{\{U\}}$ and has $d_U(x, p) \leq M$. This completes the proof of the claim. \blacksquare

Claim 3. A_n is k_n -hierarchically quasiconvex.

Proof. Claim 2 provides a constant C so that $\pi_U(\mathcal{X})$ and $\pi_U(A_n)$ C -coarsely coincide. Also, $\pi_U(A_n)$ C -coarsely coincides with $\pi_U(A'_{n-1})$ when U is not admissible. This verifies the part of the definition of hierarchical quasiconvexity governing projections; it remains to check the part governing realization points.

Consider $x \in \mathcal{X}$ so that $d_V(x, A_n) \leq t$ for each $V \in \mathfrak{S}$ and some $t \geq 0$. Let \mathcal{U} be the collection of all $U \in \mathfrak{U}_1$ so that $d_U(x, A_{n-1}) \geq 10ME$. Note that each $U \in \mathcal{U}$ is admissible. We now show that \mathcal{U} is totally orthogonal, thereby showing that \mathcal{U} is admissible.

Since each $U \in \mathcal{U}$ is \sqsubseteq -minimal, distinct elements of \mathcal{U} are \sqsubseteq -incomparable. Hence we only have to rule out the existence of $U_1, U_2 \in \mathcal{U}$ so that $U_1 \pitchfork U_2$. If such U_i existed then, up to switching U_1, U_2 , we would have $d_{U_2}(x, \rho_{U_2}^{U_1}) \leq E$ by consistency and hence $d_{U_2}(\rho_{U_2}^{U_1}, A_{n-1}) \geq 10ME - 2E > M$, contradicting the admissibility of U_1 .

Let us now define the tuple \vec{b} in the following way. For each $V \in \mathfrak{S}$, let $\pi_{V, A_{n-1}}$ be a coarse closest point projection $\mathcal{C}V \rightarrow \pi_V(A_{n-1})$. Since A_{n-1} is hierarchically quasiconvex, this map is defined in the usual way when $\mathcal{C}V$ is hyperbolic, and is either coarsely constant or coarsely the identity otherwise. We set $b_V = \pi_{V, A_{n-1}}(\pi_V(x))$ if $V \perp \mathcal{U}$. Also, we require $d_V(b_V, \rho_V^U) \leq 10E$ if $V \in \mathfrak{S}$ satisfies $U \pitchfork V$ or $U \sqsubset V$. Finally, we set $b_U = \pi_U(x)$ whenever $U \in \mathcal{U}$. It is easy to check that \vec{b} is consistent, allowing us to invoke realization (Theorem 1.12) to get a point $x' \in \mathcal{X}$. In fact, we have $x' \in B_{\mathcal{U}} \subseteq A_n$.

We now bound $d(x, x')$. In order to do so, by the uniqueness axiom, we can instead uniformly bound $d_V(x, b_V)$ for each $V \in \mathfrak{S}$. We consider the following cases.

- If $V \in \mathcal{U}$, then $b_V = \pi_V(x)$, as required.
- If $V \perp \mathcal{U}$ then $\pi_V(x)$ is within distance $t + 10EMD$ of A_{n-1} , and thus within distance $t + 10EMD$ of $b_V = \pi_{V, A_{n-1}}(\pi_V(x))$.
- Suppose that there exists $U \in \mathcal{U}$ with $U \triangleleft V$. If, by contradiction, we had $d_V(x, \rho_V^U) \geq 100ME$, then by consistency we would have $d_U(x, \rho_U^V) \leq E$. Since by definition of \mathcal{U} we have that $\pi_U(x)$ is far from $\pi_U(A_{n-1})$, we then get $d_U(A_{n-1}, \rho_U^V) > E$. By consistency applied to any $y \in A_{n-1}$, we get $\text{diam}_V(\pi_V(A_{n-1}) \cup \rho_V^U) \leq 2E$. In particular, $d_V(x, A_{n-1}) \geq 10ME$, i.e., $V \in \mathcal{U}$, a contradiction.
- Suppose that there exists $U \in \mathcal{U}$ with $U \not\sqsubseteq V$. If, by contradiction, we had $d_V(x, \rho_V^U) \geq 100ME$, then any geodesic from $\pi_V(x)$ to $\pi_{V, A_{n-1}}(x)$ stays E -far from ρ_V^U since $d_V(x, A_{n-1}) \leq 10ME$. Let $y \in A_{n-1}$ be any point so that $\pi_V(y) = \pi_{V, A_{n-1}}(x)$. By bounded geodesic image and consistency (for x and y) we have $d_U(x, A_{n-1}) \leq d_U(x, y) \leq 10E$, contradicting $U \in \mathcal{U}$.

This completes the proof of the claim. \blacksquare

Claim 4. A_n has asymptotic dimension $\leq \text{asdim } \mathcal{O}$ with control function f_n .

Proof. For $i \geq 0$, we define $A_n^{(i)} = A'_{n-1} \cup \bigcup_{\mathcal{U}} B_{\mathcal{U}}$, where the union is taken over all admissible collections \mathcal{U} of cardinality at most i . Observe that there exists I , depending only on the complexity of $(\mathcal{X}, \mathfrak{S})$, so that $A_n^{(i)} = A_n$ for all $i \geq I$. Hence it suffices to show, by induction on i , that for each $i \geq 0$, we have that $\text{asdim } A_n^{(i)} \leq \text{asdim } \mathcal{O}$, with some control function $f_n^{(i)}$.

When $i = 0$, we have $A_n^{(0)} = A'_{n-1}$, which we saw above has asymptotic dimension $\leq \text{asdim } \mathcal{O}$ with control function $f_n^{(0)} = f'_{n-1}$. Suppose for some $i \geq 0$ that we have $\text{asdim } A_n^{(i)} \leq \text{asdim } \mathcal{O}$ with control function $f_n^{(i)}$.

Write $A_n^{(i+1)} = A_n^{(i)} \cup \bigcup_{\mathcal{U}} B_{\mathcal{U}}$, where \mathcal{U} varies over all admissible totally orthogonal subsets in \mathfrak{U}_1 with $|\mathcal{U}| \leq i + 1$. Note that $\text{asdim } B_{\mathcal{U}} \leq \text{asdim } \mathcal{O}$ uniformly, since $B_{\mathcal{U}}$ is (uniformly) coarsely contained in $\prod_{U \in \mathcal{U}} \mathbf{F}_U$. By our induction hypothesis, for each $r \geq 0$, we have that $\mathcal{N}_r(A_n^{(i)})$ has asymptotic dimension $\leq \text{asdim } \mathcal{O}$. Below, we will establish the following: there exists K independent of r so that for all $r \geq 0$, and all admissible sets $\mathcal{U} \neq \mathcal{U}'$ of cardinality at most i , we have $d(B_{\mathcal{U}} - Y_r, B_{\mathcal{U}'} - Y_r) \geq r$, where $Y_r = \mathcal{N}_{Kr+K}(A_n^{(i)})$. Given this, the Union Theorem [BD01, Theorem 1] proves the claim, where the control function is $f_n = f_n^I$.

In other words, it suffices to prove that there exists $K \geq 0$ such that, if $\mathcal{U}, \mathcal{U}' \subseteq \mathfrak{U}_1$ are admissible totally orthogonal subsets of cardinality at most i , then

$$\mathcal{N}_r(B_{\mathcal{U}}) \cap \mathcal{N}_r(B_{\mathcal{U}'}) \subseteq \mathcal{N}_{Kr+K}(A'_{n-1}) \cup \mathcal{N}_{Kr+K}(B_{\mathcal{U} \cap \mathcal{U}'}).$$

Let $x \in \mathcal{N}_r(B_{\mathcal{U}}) \cap \mathcal{N}_r(B_{\mathcal{U}'})$.

Let $V \in \mathfrak{S}$. If $V \triangleleft U$ or $U \sqsubseteq V$ for some $U \in \mathcal{U} \cup \mathcal{U}'$, then $d_V(x, \rho_V^U) \leq \lambda r + \mu$, where λ, μ depend only on M and the coarse Lipschitz constants for π_V (which are independent of V). Since each $U \in \mathcal{U} \cup \mathcal{U}'$ is admissible, we then have $d_V(\rho_V^U, A'_{n-1}) \leq M$, so that $d_V(x, A'_{n-1}) \leq \lambda r + \mu + M$. If $V \perp U$ for all $U \in \mathcal{U}$ or $V \perp U$ for all $U \in \mathcal{U}'$, then $d_V(x, A'_{n-1}) \leq \lambda r + \mu + M$ by admissibility.

Suppose that V is $10(\lambda r + \mu_M)$ -relevant for $x, \mathfrak{g}_{A'_{n-1}}(x)$, i.e., $d_V(x, \mathfrak{g}_{A'_{n-1}}(x)) \geq 10(\lambda r + \mu_M)$. If $V \notin \mathcal{U}$ (respectively, \mathcal{U}'), then the preceding discussion shows that for all $U \in \mathcal{U} \cup \mathcal{U}'$, we have $V \not\triangleleft U$ and $U \not\sqsubseteq V$. On the other hand, there must exist $U \in \mathcal{U} \cup \mathcal{U}'$ with $U \perp V$, whence $V \not\sqsubseteq U$, which is impossible since $\mathcal{U} \cup \mathcal{U}' \subseteq \mathfrak{U}_1$. Hence any such relevant V lies in $\mathcal{U} \cap \mathcal{U}'$.

Let $K' = \max\{\lambda, \mu + M\}$ and let $\mathcal{U} \cap \mathcal{U}' = \{U_j\}$. We have established that:

- $d_V(x, A'_{n-1}) \leq K'r + K'$ for all V for which there exists j with $V \triangleleft U_j$ or $U_j \sqsubseteq V$, and indeed $d_V(x, \rho_V^{U_j}) \leq K'r + K'$;
- $d_V(x, A'_{n-1}) \leq K'r + K'$ for all V such that $V \perp U_j$ for all j ;

Define a tuple $\vec{t} \in \prod_{V \in \mathfrak{S}} 2^{CV}$ by $t_{U_i} = \pi_{U_i}(x)$ for all i , and so that $t_V = \rho_V^{U_i}$ if there exists i with $U_i \triangleleft V$ or $U_i \sqsubseteq V$ (for some arbitrarily-chosen i with that property if there are many – the $\rho_V^{U_i}$ all $10E$ -coarsely coincide, as can be seen by considering $\pi_V(\prod_i \mathbf{F}_{U_i})$), and $t_V = \pi_V(\mathfrak{g}_{A'_{n-1}}(x))$ otherwise.

We claim that \vec{t} is $100E$ -consistent. Indeed, let $V, W \in \mathfrak{S}$ with $V \triangleleft W$ or $V \sqsubseteq W$ or $W \sqsubseteq V$. If $t_V = \pi_V(\mathfrak{g}_{A'_{n-1}}(x))$ and $t_W = \pi_W(\mathfrak{g}_{A'_{n-1}}(x))$, then we are done by the consistency axiom. If $V = U_i$ for some i , then t_W $10E$ -coarsely coincides with $\rho_W^{U_i} = \rho_W^V$ as required. If $t_W = \pi_W(\mathfrak{g}_{A'_{n-1}}(x)) = \pi_W(g)$ and $V \triangleleft U_i$ or $U_i \sqsubseteq V$ for some U_i , then $t_V = \rho_V^{U_i}$. For each j , either $W \perp U_j$ or $W \sqsubseteq U_j$, by definition. The latter is impossible since $U_j \in \mathfrak{U}_1$, and hence $W \perp U_i$. If $V \triangleleft W$ or $W \sqsubseteq V$, it follows that $\rho_V^{U_i}$ is $10E$ -coincident with ρ_V^W . If $V \sqsubseteq W$, then $V \perp U_i$, a contradiction. Finally consider the case where there exist i, j so that $t_V = \rho_V^{U_i}$ and $t_W = \rho_W^{U_j}$. If $V \sqsubseteq W$, then $W \not\perp U_i$, so we can take $U_j = U_i$. Then the claim follows by ρ -consistency (Proposition 1.5). If $V \triangleleft W$ and $d_W(\rho_V^{U_j}, \rho_W^V) > 100E$, then $d_W(z, \rho_W^V)$ for all $z \in P_{U_j}$, whence $d_V(z, \rho_V^W) \leq E$ by consistency. But $\pi_V(z)$ $10E$ -coarsely coincides with $\rho_V^{U_i}$ since $U_i \perp U_j$.

Hence, by realization (Theorem 1.12), there exists $y \in \mathcal{X}$ with $d_V(y, t_V) \leq \theta$ for all V . If $M \geq \theta$, then $y \in B_{\{U_i\}}$. The distance formula (with fixed threshold θ , independent of r) thus provides K so that $d(x, B_{\{U_i\}}) \leq d(x, y) \leq Kr + K$, since $d_V(x, y) \leq K'r + K' + \theta$ for each V that is irrelevant for $x, \mathfrak{g}_{A'_{n-1}}(x)$ and $d_{U_i}(x, y) \leq \theta$ for each i . \blacksquare

Assertion (4) follows from Claim 1, assertion (1) follows from Claim 3, and assertion (2) follows from Claim 4. Assertion (3) holds by definition. This completes the proof. \square

5. PROOF OF THEOREM A AND COROLLARIES B AND C

In this section, we fix a uniformly locally-finite discrete geodesic space \mathcal{X} admitting a δ -relatively hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{S})$ of complexity ξ . Let \mathfrak{U}_1 be the set of \sqsubseteq -minimal elements, so that \mathcal{CU} is δ -hyperbolic for each $U \in \mathfrak{S} - \mathfrak{U}_1$. Let \mathcal{O} be the set of totally orthogonal subsets of \mathfrak{U}_1 and let $\text{asdim } \mathcal{O}$ denote the minimal uniform asymptotic dimension of $\{\prod_{V \in \mathfrak{V}} \mathbf{F}_V : \mathfrak{V} \in \mathcal{O}\}$. For $U \in \mathfrak{U}_1$, by [BHS15, Lemma 2.1] and [BD08, Theorem 32], we have that \mathcal{CU} and \mathbf{F}_U are uniformly quasi-isometric and for $\mathfrak{V} \in \mathcal{O}$ we have $\text{asdim } \prod_{V \in \mathfrak{V}} \mathcal{CV} \leq \xi \max\{\text{asdim } \mathcal{CV} : V \in \mathfrak{V}\}$. Hence $\text{asdim } \mathcal{O} \leq \xi n$, where n is the minimal uniform asymptotic dimension of $\{\mathcal{CU} : U \in \mathfrak{U}_1\}$.

Remark. In order to apply Corollary 3.3, we must assume that for each $U \in \mathfrak{S}$ with \mathcal{CU} a δ -hyperbolic space, the projection π_U is uniformly coarsely surjective. By the proof of Proposition 1.16 of [DHS16], we can always assume that this holds (see also Remark 1.3 of [BHS15]). Hence, in the proof of Theorem 5.2, we can make this coarse surjectivity assumption and thus apply Corollary 3.3.

Definition 5.1 (Level, P_ℓ , Δ_ℓ). Define the *level* of $U \in \mathfrak{S}$ to be 1 if U is \sqsubseteq -minimal and inductively define the level of $U \in \mathfrak{S}$ to be ℓ if $\ell - 1$ is the maximal integer such that there exists $V \in \mathfrak{S}$ of level $\ell - 1$ with $V \sqsubseteq U$. Let P_ℓ be the maximal cardinality of pairwise-orthogonal sets in \mathfrak{S} each of whose elements has level ℓ . Let Δ_ℓ be the maximal uniform asymptotic dimension of \mathcal{CU} with U of level ℓ and \mathcal{CU} hyperbolic.

Theorem 5.2. *Let \mathcal{X} be a uniformly locally-finite discrete geodesic space admitting a δ -relatively HHS structure $(\mathcal{X}, \mathfrak{S})$ and let ξ, n , and the Δ_ℓ be as above. Assume that $n < \infty$.*

Then $\text{asdim } \mathcal{X} \leq n\xi + \sum_{\ell=2}^{\xi} P_{\ell}\Delta_{\ell} < \infty$. In particular, if $(\mathcal{X}, \mathfrak{S})$ is an HHS, then $\text{asdim } \mathcal{X} \leq \sum_{\ell=1}^{\xi} P_{\ell}\Delta_{\ell} < \infty$.

Observe that when $(\mathcal{X}, \mathfrak{S})$ is actually an HHS, Corollary 3.3 automatically gives $n < \infty$.

Before proving Theorem 5.2, we record a lemma whose proof is an immediate consequence of [CdIH14, Lemma 3.B.6].

Lemma 5.3. *Let \mathcal{X} be a (not necessarily locally-finite) (r_0, r_1) -discrete geodesic space with a δ -relatively hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{S})$. Let $\mathfrak{U} \subseteq \mathfrak{S}$ be closed under nesting. Then $\widehat{\mathcal{X}}_{\mathfrak{U}}$ is quasi-isometric to a connected graph Γ , with constants independent of \mathfrak{U} .*

Proof of Theorem 5.2. By Proposition 4.1, for $R \geq 0, x \in \widehat{\mathcal{X}}_{\mathfrak{U}_1}$, we have $\text{asdim } \psi^{-1}(B^{\widehat{\mathcal{X}}_{\mathfrak{U}_1}}(x, R)) \leq \text{asdim } \mathcal{O}$ uniformly, where $\psi: \mathcal{X} \rightarrow \widehat{\mathcal{X}}_{\mathfrak{U}_1}$ is the Lipschitz map provided by Proposition 2.2. The space \mathcal{X} is a geodesic space, by Lemma 5.3. Hence, we may apply Theorem 1.1, which yields $\text{asdim } \mathcal{X} \leq \text{asdim } \mathcal{O} + \text{asdim } \widehat{\mathcal{X}}_{\mathfrak{U}_1}$.

Now, by Proposition 2.4, $(\widehat{\mathcal{X}}_{\mathfrak{U}_1}, \mathfrak{S} - \mathfrak{U}_1)$ is an HHS of complexity $\xi - 1$, and $\widehat{\mathcal{X}}_{\mathfrak{U}_1}$ is uniformly quasi-isometric to a geodesic space (a graph) by Lemma 5.3. Observe that for each \sqsubseteq -minimal $U \in \mathfrak{S} - \mathfrak{U}_1$, we have that the associated subspace $\mathbf{F}_U \subset \widehat{\mathcal{X}}_{\mathfrak{U}_1}$ is uniformly quasi-isometric to $\mathcal{C}U$, and thus, by induction,

$$\text{asdim } \mathcal{X} \leq n\xi + \sum_{\ell=2}^{\xi} P_{\ell}\Delta_{\ell},$$

which is finite since $P_{\ell} \leq \xi$ for all ℓ by [BHS15, Lemma 2.1] and $\Delta_{\ell} < \infty$ by Corollary 3.3. This yields the desired bound for $(\mathcal{X}, \mathfrak{S})$ an HHS, since then we may take $n = \Delta_1$. \square

In the case of the mapping class group, sharper bounds on the asymptotic dimensions of curve graphs are known. Webb has a combinatorial argument which gives a bound which is exponential in the complexity [Web15]. We will make use of a much tighter bound due to Bestvina–Bromberg [BB15]. Using this we will now prove Corollary B:

Proof of Corollary B. We use the hierarchically hyperbolic structure on $\mathcal{MCG}(S)$ from [BHS15, Section 11], where S is a surface of complexity $\xi(S)$, and \mathfrak{S} is the set of essential subsurfaces up to isotopy, and for each $U \in \mathfrak{S}$, the space $\mathcal{C}U$ is the curve graph. Let L be the maximum level of a subsurface, i.e., the maximal ℓ so that \mathfrak{S} has a chain $U_1 \sqsubset U_2 \sqsubset \dots \sqsubset U_{\ell} = S$.

Let $\{U_1, \dots, U_k\}$ be a collection of pairwise-disjoint subsurfaces, each of level exactly $\ell > 1$. Then each U_i contains a subsurface U'_i of level exactly $\ell - 1$, and the complement in U_i of U'_i has level 1 unless it is a degenerate subsurface. Hence U_i contains at least $\ell - 1$ disjoint subsurfaces of level 1, so $(\ell - 1)P_{\ell} \leq L$ for all $\ell \geq 2$ while $P_1 = L = \xi(S)$ (note, U_i contains at most ℓ disjoint subsurfaces of level 1). As shown in [BB15], $\text{asdim } \mathcal{C}U \leq 2\ell + 3$ uniformly. Thus $\Delta_{\ell} \leq 2\ell + 3$, and Theorem 5.2 gives:

$$\text{asdim}(\mathcal{MCG}(S)) \leq 5L + L \sum_{\ell=2}^L \frac{(2\ell + 3)}{\ell - 1} \leq 2L^2 + 3L \log L + 8L.$$

Observing that $L \leq \xi(S)$ provided $\xi(S) \geq 2$ completes the proof. \square

We now prove Corollary C:

Proof of Corollary C. As noted in [BHS14], Teichmüller space with either of the two metrics mentioned is a hierarchically hyperbolic space; for details see the corresponding discussion for the mapping class group in [BHS15, Section 11] which applies *mutatis mutandis* in the present context. The index-set \mathfrak{S} consists of all isotopy classes of essential subsurfaces (only non-annular ones in the case of the Weil–Petersson metric). For each $U \in \mathfrak{S}$ which is not an

annulus, \mathcal{CU} is the curve graph of U ; For annular U the space \mathcal{CU} is a combinatorial horoball over the annular curve graph. Observing that $\text{asdim } \mathcal{CU} \leq 2$ when \mathcal{CU} is the horoball over an annular curve graph, the claim now follows as in the proof of Corollary B, albeit with an extra additive term of $\xi(S)$ since the lowest complexity terms have asymptotic dimension 2 in the present case instead of 1. (Since Weil–Petersson doesn't have contributions from annuli at all, one obtains a sharper bound than we record here.) \square

6. QUOTIENTS OF HHG

In this section, we study quotients of hierarchically hyperbolic groups. It will be a standing assumption throughout this section that H is a *hierarchically embedded subgroup* of the HHG (G, \mathfrak{S}) , as defined by the following, cf. [DGO11]:

Definition 6.1 (Hierarchically hyperbolic embedded subgroup). Let (G, \mathfrak{S}) be an HHG, let $S \in \mathfrak{S}$ be \sqsubseteq -maximal and let $H \leq G$. Given a (possibly infinite) generating set \mathcal{T} of G , then we write $H \hookrightarrow_{hh} (G, \mathfrak{S})$ if the following hold:

- \mathcal{CS} is the Cayley graph of G with respect to \mathcal{T} and π_S is the inclusion;
- $\mathcal{T} \cap H$ generates H ;
- H is hyperbolic embedded in (G, \mathcal{T}) . Recall that this means that $\text{Cay}(G, H \cup \mathcal{T})$ is hyperbolic and H is proper with respect to the metric \hat{d} obtained from measuring the length of a shortest path $\gamma \subset \text{Cay}(G, H \cup \mathcal{T})$ with the property that between pairs of vertices in $H \cap \gamma$ the only edges allowed are those from \mathcal{T} .

Throughout this section we let N denote a subgroup $N \triangleleft H \leq G$, and we let \hat{N} denote its normal closure in G . When dealing with different HHS structures $\mathfrak{S}, \mathfrak{T}$, and when it is necessary to distinguish between the two, we will use the notation $\mathcal{C}_{\mathfrak{S}}U, \mathcal{C}_{\mathfrak{T}}V$ for $U \in \mathfrak{S}, V \in \mathfrak{T}$ instead of $\mathcal{CU}, \mathcal{CV}$. Our main result is:

Theorem 6.2. *Let (G, \mathfrak{S}) be an HHG and let $H \hookrightarrow_{hh} (G, \mathfrak{S})$. Then there exists a finite set $F \subset H - \{1\}$ such that for all $N \triangleleft H$ with $F \cap N = \emptyset$, the group G/\hat{N} admits a relative HHG structure $(G/\hat{N}, \mathfrak{S}_N)$ where:*

- $\mathfrak{S}_N = (\mathfrak{S} \cup \{gH\}_{g \in G})/\hat{N}$ and hence $\xi(\mathfrak{S}_N) \leq \max\{\xi(\mathfrak{S}), 2\}$;
- $\mathcal{C}_{\mathfrak{S}_N}S_N = \text{Cay}(G/\hat{N}, \mathcal{T}/\hat{N} \cup H/N)$, where $S_N \in \mathfrak{S}_N$ is \sqsubseteq -maximal and \mathcal{T} is as in Definition 6.1;
- for each $U \in \mathfrak{S}_N - \{S_N\}$, either $\mathcal{C}_{\mathfrak{S}_N}U$ is isometric to $\mathcal{C}_{\mathfrak{S}}U'$ for some $U' \in \mathfrak{S}$, or $\mathcal{C}_{\mathfrak{S}_N}U$ is isometric to a Cayley graph of H/N .

In particular, if N avoids F and H/N is hyperbolic, then G/\hat{N} is an HHG.

We postpone the proof until after explaining the necessary tools. The definition of the index set \mathfrak{S}_N in the above theorem means the following: by the definition of a hierarchically hyperbolic group structure, G , and hence $\hat{N} \leq G$, acts on \mathfrak{S} — see Section 1.2.3. Also, \hat{N} acts on the set $\{gH\}_{g \in G}$ in the obvious way. Thus \hat{N} acts on the set $\mathfrak{S} \cup \{gH\}_{g \in G}$, and \mathfrak{S}_N is the quotient of $\mathfrak{S} \cup \{gH\}_{g \in G}$ by this action.

Remark 6.3 (Geometric separation and quasiconvexity). Since $H \hookrightarrow_h (G, \mathcal{T})$, we have that H is *geometrically separated* (with respect to \mathcal{CS}), i.e., for each $\epsilon \geq 0$ there exists $M(\epsilon)$ so that, if $\text{diam}_{\mathcal{CS}}(H \cap \mathcal{N}_{\epsilon}(gH)) \geq M(\epsilon)$, then $g \in H$. Indeed, for any $\epsilon > 0$, if we can choose g so that $\text{diam}_{\mathcal{CS}}(H \cap \mathcal{N}_{\epsilon}(gH))$ is arbitrarily large, then H contains infinitely many elements h' with $\hat{d}(1, h') \leq 2\epsilon + 1$, which is impossible. In fact, [Sis12, Theorem 6.4] shows that $\text{Cay}(G, \mathcal{T})$ is (metrically) hyperbolic relative to $\{gH\}$, and in particular there exists M such that $H \subset \text{Cay}(G, \mathcal{T})$ is M -quasiconvex, by [DS05, Lemma 4.3].

Remark 6.4 (H is hyperbolic). In our situation, the hyperbolically embedded subgroup $H \hookrightarrow_h (G, \mathcal{T})$ must be hyperbolic. This holds since (G, \mathcal{T}) is already hyperbolic (even before adding H to the generating set); thus, hyperbolicity of H follows from the fact that $H \subset (G, \mathcal{T})$ is quasiconvex and H acts properly (indeed, the word-metric $d_{\mathcal{T}}$ restricted to H is bounded below by the auxiliary metric on H from Definition 6.1, which is proper). This hyperbolicity is not a significant restriction on the subgroup, since it holds in the cases of interest, including the case where G is the mapping class group of a surface and $H \cong \mathbb{Z}$ is generated by a pseudo-Anosov element.

6.1. Pyramid spaces. In this section we define the *pyramid spaces* which are hyperbolic spaces, associated to G and G/\widehat{N} . We also describe a new hierarchically hyperbolic structure on G in which the pyramid space associated to G replaces $\text{Cay}(G, \mathcal{T})$. We begin by recalling from [GM08] the notion of a combinatorial horoball, and the attendant hyperbolic cone construction from [DGO11].

Definition 6.5 (Combinatorial horoball, hyperbolic cone). Let Γ be a graph. The *combinatorial horoball* $\mathcal{H}(\Gamma)$ is the graph formed from $\Gamma \times (\mathbb{N} \cup \{0\})$ by adding the following edges: for each $r \in \mathbb{N} \cup \{0\}$ and each vertex $v \in \Gamma$, join $(v, r), (v, r+1)$ by an edge. For each r and each $v, v' \in \Gamma$, join $(v, r), (v', r)$ if $d_{\Gamma}(v, v') \in (0, 2^r]$.

Given $r \in \mathbb{N}$, the *hyperbolic cone* $\mathfrak{C}(\Gamma, r)$ of *radius* r over Γ is obtained from $\mathcal{H}(\Gamma)$ by adding a vertex v , called the *apex* of the cone, and joining v to each vertex (w, s) of $\mathcal{H}(\Gamma)$ for which $s \geq r$. We endow $\mathfrak{C}(\Gamma, r)$ with the usual graph metric.

Lemma 6.43 of [DGO11] says that for any choice of Γ and any $r \in \mathbb{N}$, the graph $\mathfrak{C}(\Gamma, r)$ is δ -hyperbolic, where δ may be chosen independently of Γ and r .

Definition 6.6 (Pyramid spaces $\text{Pyr}(G)_r$ and $\text{Pyr}(G/\widehat{N})_r$). For $r \geq 1$, the *pyramid space* $\text{Pyr}(G)_r$ associated to (G, \mathcal{T}) and H is obtained from $\text{Cay}(G, \mathcal{T})$ by attaching the hyperbolic cone of radius r over each coset gH , with apex v_{gH} . Let d_{Δ} be the graph metric on $\text{Pyr}(G)_r$. Likewise, $\text{Pyr}(G/\widehat{N})_r$ is obtained from $\text{Cay}(G/\widehat{N}, \mathcal{T}/\widehat{N})$ by attaching a radius- r hyperbolic cone over each coset of H/N .

Proposition 6.7. *There exists $\delta \geq 1$ with the following properties:*

- (1) $\text{Pyr}(G)_r$ is δ -hyperbolic for all sufficiently large r ;
- (2) for each sufficiently large $r > 0$, there exists a finite set $F_r \subset H - \{1\}$ such that if $N \cap F_r = \emptyset$, then $\text{Pyr}(G/\widehat{N})_r$ is δ -hyperbolic.

Proof. Assertion (1) holds by [DGO11, Lemma 6.45]. To prove assertion (2), consider the action of G on the δ -hyperbolic space $\text{Pyr}(G)_r$. By construction, the set $\{v_{gH}\}$ of apices is G -invariant, $2r$ -separated, and each v_{gH} is fixed by the *rotation subgroup* H^g , i.e., $(\{v_{gH}\}, \{H^g\})$ is a $2r$ -separated *rotating family* in the sense of [DGO11, Definition 5.1]. Observe that [DGO11, Corollary 6.36] provides a finite set $F_r \subset H - \{1\}$ such that $\{N^g\}$ is a $\frac{2r}{\delta}$ -rotating family, with respect to the action of $\text{Pyr}(G)_r$, provided $N \triangleleft H$ avoids F_r . (This means that $(\{v_{gH}\}, \{N^g\})$ is a $2r$ -separated rotating family that also satisfies the *very rotating condition* of [DGO11, Definition 5.1].) Proposition 5.28 of [DGO11] provides r_0 so that $\text{Pyr}(G/\widehat{N})_r$ is uniformly hyperbolic when $r \geq r_0$. Enlarging δ completes the proof. \square

Lemma 6.8. *There exist Q, r_0 so that $\mathfrak{C}(H, r) \subset \text{Pyr}(G)_r$ is Q -quasiconvex if $r \geq r_0$.*

Proof. Choose r sufficiently large so that $\text{Pyr}(G)_r$ is δ -hyperbolic, using Proposition 6.7, and let \mathfrak{C} be the hyperbolic cone of radius r over H .

Define a map $\mathfrak{l}: \text{Pyr}(G)_r \rightarrow \mathfrak{C}$ as follows. First, for each $x \in \mathfrak{C}$, let $\mathfrak{l}(x) = x$. Next, for each $g \in G$, let $\mathfrak{l}(g)$ be some $h \in H$ so that $d_{\mathcal{CS}}(g, H) = d_{\mathcal{CS}}(g, h)$. (This is coarsely unique since H is M -quasiconvex in \mathcal{CS} by Remark 6.3.) If $gH \neq H$ and $(y, n) \in \mathfrak{C}(gH, r)$ is not the apex

v_{gH} , then let $\mathfrak{l}(y, n) = \mathfrak{l}(y)$. By Claim 1, \mathfrak{l} sends gH to a uniformly bounded subset of H , and we define $\mathfrak{l}(v_{gH})$ to be an arbitrarily-chosen point in $\mathfrak{l}(gH)$.

Claim 1. There exists a constant K_1 so that $\text{diam}_{\mathcal{CS}}(\mathfrak{l}(gH)) \leq K_1$ for all $gH \neq H$.

Proof of Claim 1. Apply M -quasiconvexity and separation of gH, H (Remark 6.3). \blacksquare

Claim 2. There exists a constant K_2 so that $\text{d}_{\mathcal{CS}}(\mathfrak{l}(x), \mathfrak{l}(y)) \leq K_2 \text{d}_{\mathcal{CS}}(x, y)$ for all $x, y \in G$.

Proof of Claim 2. Apply quasiconvexity and the definition of \mathfrak{l} . \blacksquare

Claim 3. There exists K_3 so that, if r is sufficiently large, then $\text{d}_{\Delta}(\mathfrak{l}(x), \mathfrak{l}(y)) \leq K_3 \text{d}_{\Delta}(x, y) + K_3$ for all $x, y \in \text{Pyr}(G)_r$.

Proof of Claim 3. Let $[x, y]_{\Delta}$ be a $\text{Pyr}(G)_r$ -geodesic from $x \in G$ to $y \in G$, so that $[x, y]_{\Delta} = \alpha_1 \beta_1 \cdots \alpha_n \beta_n \alpha_{n+1}$, where each β_i lies in some hyperbolic cone $\mathfrak{C}(g_i H, r)$ and each α_i avoids $v_{gH} \cup (gH \times [1, \infty))$ for each coset gH (i.e., α_i is a possibly trivial path in \mathcal{CS}).

For each i , let $\tilde{\beta}_i$ be a geodesic in \mathcal{CS} joining the endpoints a_i, b_i of β_i (which lie in $g_i H$). By Claim 1, we have $\text{d}_{\Delta}(\mathfrak{l}(a_i), \mathfrak{l}(b_i)) \leq \text{d}_{\mathcal{CS}}(\mathfrak{l}(a_i), \mathfrak{l}(b_i)) \leq K_1$.

By Claim 2, we have $\text{d}_{\mathcal{CS}}(\mathfrak{l}(b_{i-1}), \mathfrak{l}(a_i)) \leq K_2 \text{d}_{\mathcal{CS}}(b_{i-1}, a_i)$. Now, since α_i is a $\text{Pyr}(G)$ -geodesic, we have $\text{d}_{\mathcal{CS}}(b_{i-1}, a_i) = \text{d}_{\Delta}(b_{i-1}, a_i)$, whence $\text{d}_{\mathcal{CS}}(\mathfrak{l}(b_{i-1}), \mathfrak{l}(a_i)) \leq K_2 \text{d}_{\Delta}(b_{i-1}, a_i)$. Finally, $\text{d}_{\mathcal{CS}}(\mathfrak{l}(b_{i-1}), \mathfrak{l}(a_i)) \geq \text{d}_{\Delta}(\mathfrak{l}(b_{i-1}), \mathfrak{l}(a_i))$, so $\text{d}_{\Delta}(\mathfrak{l}(b_{i-1}), \mathfrak{l}(a_i)) \leq K_2 \text{d}_{\Delta}(b_{i-1}, a_i)$. Hence $\text{d}_{\Delta}(\mathfrak{l}(x), \mathfrak{l}(y)) \leq K_2 \sum_i |\alpha_i| + K_1 n \leq \max\{K_1, K_2\} \text{d}_{\Delta}(x, y)$, and we are done.

The other possibilities are that $x \in \mathfrak{C}(g_0 H), y \in \mathfrak{C}(g_{n+1} H)$ or $x \in \mathfrak{C}(g_0 H), y \in G$, so that we consider paths of the form $\beta_0 \alpha_1 \beta_1 \cdots \alpha_n \beta_n \alpha_{n+1}$ or $\beta_0 \alpha_1 \beta_1 \cdots \alpha_n \beta_n \alpha_{n+1} \beta_{n+1}$. We now argue as above and conclude that $\text{d}_{\Delta}(\mathfrak{l}(x), \mathfrak{l}(y)) \leq 2K_1 + \max\{K_1, K_2\} \text{d}_{\Delta}(x, y)$. \blacksquare

The lemma follows easily from Claim 3. \square

Remark 6.9 (Choosing r). Fix r as in Proposition 6.7 and assume $r \geq 10^9 \delta EQ$, where δ is as in Proposition 6.7 and Q as in Lemma 6.8. Let $\text{Pyr}(G) = \text{Pyr}(G)_r$ and, for each gH , let $\mathfrak{C}(gH) = \mathfrak{C}(gH, r) \subset \text{Pyr}(G)$. Later, we will impose additional assumptions on the size of r .

Definition 6.10 (Push-off). Let α be a geodesic in $\text{Pyr}(G)$ with endpoints in G . The path γ in $\mathcal{CS} = \text{Cay}(G, \mathcal{T})$ is a *push-off* of α if it is obtained by replacing sub-geodesics in hyperbolic cones with geodesics in \mathcal{CS} ; we call these new subpaths *replacement paths*.

Lemma 6.11. *Let Z be a δ -hyperbolic space and let \mathcal{Q} be a collection of Q -quasiconvex subspaces. Let α be a geodesic of Z and let $H \subseteq Z$ be the union of α and every $Y \in \mathcal{Q}$ with $\alpha \cap Y \neq \emptyset$. Then H is $(Q + 2\delta)$ -quasiconvex.*

Proof. Let γ be a geodesic joining points $h, h' \in H$. If $h, h' \in \alpha$, then we are done since $\gamma \subset \mathcal{N}_{2\delta}(\alpha)$. Next suppose $h \in Y, h' \in Y'$ for some $Y, Y' \in \mathcal{Q}$. Consider closest points p, p' of $Y \cap \alpha, Y' \cap \alpha$ to h, h' respectively, and consider the geodesic quadrilateral $\gamma \rho' \beta \rho^{-1}$, where β is the subgeodesic of α from p' to p and ρ', ρ join h, h' to p, p' . Then every point of γ lies either 2δ -close to a point of β or lies 2δ -close to a point of $\rho \cup \rho'$ and hence $(Q + 2\delta)$ -close to a point of $Y \cup Y'$. A virtually identical argument works if $h \in \alpha$ and $h' \in Y' \in \mathcal{Q}$. \square

Lemma 6.12. *There exists $C = C(\delta, Q)$ so that, if r is sufficiently large, then the following holds for all $x, y \in \text{Pyr}(G)$ and all gH : if $\text{d}_{\Delta}(x, \mathfrak{C}(gH)), \text{d}_{\Delta}(y, \mathfrak{C}(gH)) \leq 2\delta + Q$, and $\text{d}_{\Delta}(x, y) \geq C$, then any geodesic from x to y intersects $\mathfrak{C}(gH)$ in an interior point of $[x, y]$.*

Proof. Let $x', y' \in \mathfrak{C}(gH)$ satisfy $\text{d}_{\Delta}(x, x'), \text{d}_{\Delta}(y, y') \leq 2\delta + Q$ and let γ, γ' be geodesics joining x, y and x', y' respectively. Let α, β be geodesics joining x, x' and y, y' respectively. Examining the quadrilateral $\alpha \gamma' \beta^{-1} \gamma^{-1}$ provides $a, b \in \gamma$ so that $\text{d}_{\Delta}(a, x), \text{d}_{\Delta}(b, y) \leq 10\delta$, while $\text{d}_{\Delta}(a, \mathfrak{C}(gH)), \text{d}_{\Delta}(b, \mathfrak{C}(gH)) \leq 2\delta + Q$ (by Lemma 6.8), and $C \geq \text{d}_{\Delta}(a, b) \geq C - 20\delta$. Suppose that the subpath $[a, b]$ of γ from a to b does not enter $\mathfrak{C}(gH)$ (except possibly at

the endpoint). Then this path projects to a path of length at most $KC + K$, for $K = K(Q)$, that lies in gH and joins points a', b' respectively at distance $\leq 2\delta + Q$ from a, b . Thus, if r is sufficiently large, we have a path of length $2\log_2(KC + K) + 1$ in $\mathfrak{C}(gH)$ joining a', b' , whence $d_\Delta(a, b) \leq 4\delta + Q + 2\log_2(KC + K) + 1$. On the other hand, $d_\Delta(a, b) \geq C - 20\delta$, and this is a contradiction provided C was chosen sufficiently large (in terms of Q and δ). \square

Lemma 6.13 (Push-offs are close to \mathcal{CS} -geodesics). *Suppose that r is sufficiently large, in terms of M, δ . Then there exists D , independent of r , so that any push-off of a geodesic in $\text{Pyr}(G)$ that starts and ends in G is a (D, D) -quasigeodesic and lies within Hausdorff distance D (in \mathcal{CS}) from a geodesic in \mathcal{CS} with the same endpoints.*

Proof. Let $[a, b] = \alpha_0\beta_1 \cdots \alpha_{n-1}\beta_n\alpha_{n+1}$ be a $\text{Pyr}(G)$ -geodesic with $a, b \in \mathcal{CS}$, where each α_i is a \mathcal{CS} -geodesic and each β_i lies in some $\mathfrak{C}(g_iH)$. For each i , let $\hat{\beta}_i$ be a \mathcal{CS} -geodesic joining the endpoints of β_i , so that $\gamma = \alpha_0\hat{\beta}_1 \cdots \alpha_{n-1}\hat{\beta}_n\alpha_{n+1}$ is a push-off of $[a, b]$. Let $H = [a, b] \cup_i \mathfrak{C}(g_iH)$, so that H is $(Q + 2\delta)$ -quasiconvex by Lemma 6.11. Let $p: \text{Pyr}(G) \rightarrow H$ be the projection, and fix a 1-Lipschitz parameterization $\gamma: I \rightarrow \mathcal{CS}$ of γ .

Assuming r is large enough, we now define a coarsely Lipschitz projection $q: \mathcal{CS} \rightarrow I$, with constants bounded in terms of M, δ . The existence of such map q easily implies that γ is a quasigeodesic with constants depending on M, δ only, and hence that it is also Hausdorff close to a geodesic, as required. Write $I = A_0 \cup B_1 \cup \cdots \cup A_{n-1} \cup B_n \cup A_{n+1}$, where $\gamma|_{A_i} = \alpha_i$ and $\gamma|_{B_i} = \hat{\beta}_i$. Given $g \in \mathcal{CS}$, if $p(g) \in \alpha_i$, then let $q(g)$ be chosen in A_i so that $\gamma(q(g)) = p(g)$. Otherwise, $p(g) \in g_iH$ for some i , and we let $q(g)$ be chosen in B_i so that $\hat{\beta}_i(q(g))$ is the closest-point projection of $p(g)$ on $\hat{\beta}_i$, i.e., $d_{\mathcal{CS}}(p(g), \hat{\beta}_i(q(g))) = d_{\mathcal{CS}}(p(g), \hat{\beta}_i)$.

Our goal is now to show that q is coarsely Lipschitz whenever r is large enough. It suffices to bound $|q(g) - q(h)|$ for $g, h \in \mathcal{CS}$ satisfying $d_{\mathcal{CS}}(g, h) \leq 1$.

Let $g, h \in \mathcal{CS}$ satisfy $d_{\mathcal{CS}}(g, h) \leq 1$. Then there exists $K = K(Q, \delta)$ so that $d_\Delta(p(g), p(h)) \leq K$. First of all, we show that we can bound $|q(g) - q(h)|$ whenever r is large enough in any of the following cases:

- $p(g), p(h)$ both belong to some g_iH ,
- $p(g), p(h)$ each lie on some α_i .

In fact, in the first case we can use the fact that, provided r is much larger than K , we have $d_{\mathcal{CS}}(p(g), p(h)) \leq K' = K'(K)$, combined with the fact that the closest point projection on $\hat{\beta}_i$ is coarsely Lipschitz. In the second case, the fact that $p(g), p(h)$ are connected by a subgeodesic of $[a, b]$ of length at most K again ensures that $d_{\mathcal{CS}}(p(g), p(h)) \leq K' = K'(K)$ whenever r is large enough.

Up to switching g, h , there is only one case left to analyze: Suppose that there exists i so that $p(g) \in g_iH, p(h) \notin g_iH$; then $p(h) \in g_jH \cup \alpha_j$ for some j . Let $a' \in g_iH$ be the entrance point in g_iH of the subpath of $[a, b]$ joining α_j to g_iH . Then we claim that there exists $C' = C'(\delta, Q)$ so that $d_{\mathcal{CS}}(a', p(g)) \leq C'$. Since a similar statement holds for h as well if $p(h) \in g_jH$, following arguments similar to the ones above we can then easily get the required bound on $|q(g) - q(h)|$ provided r is much larger than K and C' .

Consider a geodesic quadrilateral determined by $p(g), p(h), a', c$, in that order, where $c \in [a, b]$ and $[c, a']$ intersects $\mathfrak{C}(g_iH)$ only in a' , and either $c, p(h) \in g_jH$ or $c = p(h)$. Choose s (independent of r) so that $\text{diam}(\mathcal{N}_Q(\mathfrak{C}(gH)) \cap \mathcal{N}_{2\delta+2Q}(\mathfrak{C}(g'H))) \leq s$ whenever $gH \neq g'H$. Suppose by contradiction that $d_\Delta(a', p(g)) > 10\delta + s + C + K$, for C as in Lemma 6.12. Choose $x, y \in [p(g), a']$ with $d_\Delta(p(g), x) = 3\delta + K$ and $d_\Delta(p(g), y) = 5\delta + K + s$. Then, since $d_\Delta(x, [p(g), p(h)]), d_\Delta(y, [p(g), p(h)]) \geq K + 3\delta - K > 2\delta$, we have that x, y are 2δ -close to $[a', c]$ or $[p(h), c]$. However, the former is ruled out by Lemma 6.12 since $d_\Delta(a', x), d_\Delta(a', y) > C + 2\delta$ and the fact that $[a', c]$ does not have interior points in $\mathfrak{C}(g_iH)$. Hence we must have $p(h) \neq c$, so that $p(h), c \in g_jH$, and $d_\Delta(x, [p(h), c]), d_\Delta(y, [p(h), c]) \leq 2\delta$, which is impossible,

in view of the definition of s , since $d_\Delta(x, g_i H), d_\Delta(y, g_i H) \leq Q$ and $d_\Delta(x, y) > s$. Hence, we showed $d_{CS}(a', p(g)) \leq C'$ for $C' = 10\delta + s + C + K$, as required. \square

6.1.1. *An alternative HHG structure on G .* It will be convenient to add the cosets of H to the HHG structure of G and, to do so, we also must replace CS with $\text{Pyr}(G)$. The index set of the new structure will include \mathfrak{S} as a proper subset. For each element of $W \in \mathfrak{S} - \{S\}$, the associated hyperbolic space CW is the same in both structures; the hyperbolic space associated to S will differ, so we denote the two spaces by $C_{\mathfrak{S}}S$ and $C_{\mathfrak{T}}S$; sometimes for emphasis we will, more generally, use the notation $C_{\mathfrak{S}}W$ and $C_{\mathfrak{T}}W$ to emphasize which structure we are considering at the time.

Proposition 6.14. *The following is an HHG structure on G :*

- the index set, \mathfrak{T} , contains all the elements of \mathfrak{S} together with one element for each coset $\{gH\}_{g \in G}$;
- \sqsubseteq and \perp restricted to \mathfrak{S} are unchanged, each gH is only nested into S and not orthogonal to anything;
- $C_{\mathfrak{T}}S$ is $\text{Pyr}(G)$, while $C_{\mathfrak{T}}U = C_{\mathfrak{S}}U$ for $U \in \mathfrak{S} - \{S\}$. Set $C_{\mathfrak{T}}gH = \text{Cay}(gH, \mathcal{T} \cap H)$;
- ρ_V^U is unchanged for $U, V \in \mathfrak{S} - \{S\}$ (when defined);
- for $U \in \mathfrak{S}$, the map $\rho_U^S: \text{Pyr}(G) \rightarrow CU$ is unchanged on $CS \subset \text{Pyr}(G)$, while $\rho_U^S((gh, s)) = \rho_U^S(gh)$ for each $g \in G, h \in H, s < \infty$ and $\rho_U^S(v_{gH}) = \cup_{h \in H} \rho_U^S(gH)$;
- $\rho_{gH}^S(x)$ is the set of entrance points in $\mathfrak{C}(gH)$ of all geodesics $[x, v_{gH}]$, when $x \notin \mathfrak{C}(gH)$, and otherwise $\rho_{gH}^S(x) = gH$, while $\rho_S^{gH} = v_{gH}$;
- ρ_{gH}^U is $\rho_{gH}^S(\rho_S^U)$ for $U \in \mathfrak{S} - \{S\}$, while $\rho_U^{gH} = \pi_U(gH)$ for $U \in \mathfrak{S} - \{S\}$;
- π_U is unchanged for $U \in \mathfrak{S}$ and is the composition $\rho_{gH}^S \circ \pi_S$ for each gH .

Before proving the proposition, we record two lemmas:

Lemma 6.15. *There exists $C \geq 1$, independent of r , so that for each $x, g \in G$, the set of entry points of $\text{Pyr}(G)$ -geodesics $[x, v_{gH}]$ in $\mathfrak{C}(gH)$ is within Hausdorff distance C of $\{x' \in gH : d_{CS}(x, x') = d_{CS}(x, gH)\}$. Hence there exists C with $\text{diam}(\rho_{gH}^S(a)) \leq C$ for all $a, g \in G$.*

Proof. This follows from M -quasiconvexity of gH in CS (Remark 6.3) and Lemma 6.13. \square

Lemma 6.16. *There exists $C \geq 1$, independent of r , so that:*

- (1) $\text{diam}_{CU}(\pi_U(gH)) \leq C$ for all $g \in G$ for all $U \in \mathfrak{S} - \{S\}$;
- (2) for all $g'H \neq gH$, $x, y \in g'H$, and all geodesics $[x, v_{gH}], [y, v_{gH}]$ of $\text{Pyr}(G)$ from x, y to v_{gH} , the entry points a_x, a_y of $[x, v_{gH}], [y, v_{gH}]$ in $\mathfrak{C}(gH)$ satisfy $d_{gH}(a_x, a_y) \leq C$.

Proof. The first assertion follows from bounded geodesic image and the fact that H acts properly on $\text{Cay}(G, \mathcal{T})$. The second follows from Lemma 6.15 and Remark 6.3. \square

We can now prove the proposition:

Proof of Proposition 6.14. All aspects of Definition 1.4 involving only the $\sqsubseteq, \perp, \pitchfork$ relations (but not the projections) are obviously satisfied. Note that G acts cofinitely on \mathfrak{T} and each $g \in G$ induces an isometry $C_{\mathfrak{T}}U \rightarrow C_{\mathfrak{T}}(gU)$ for each $U \in \mathfrak{T}$. Abusing notation slightly, we denote by gH the subgraph of $\text{Cay}(G, \mathcal{T})$ spanned by the vertices of gH , which is connected since $\mathcal{T} \cap H$ generates H , and which is hyperbolic by Remark 6.4.

Projections π_U are well-defined and coarsely Lipschitz: This is automatic for $U \in \mathfrak{S}$. For each gH and $a \in G$, the projection $\pi_{gH}(a) = \rho_{gH}^S(\pi_S(a)) = \rho_{gH}^S(a)$ is bounded by Lemma 6.15. We now verify that π_{gH} is coarsely Lipschitz; it suffices to verify that ρ_{gH}^S is coarsely Lipschitz on $G \subset CS$. Let γ be a geodesic in $\text{Pyr}(G)$ joining $a, b \in G$, and let α be

a push-off of γ , so that $d_\Delta(a, b) = |\gamma| \leq |\alpha|$. By Lemma 6.13, α lies at Hausdorff distance at most D from a geodesic α' of \mathcal{CS} , and the claim follows.

Lemma 6.16 says that $\pi_U(gH)$ is bounded for $g \in G, U \in \mathfrak{S}$, so ρ_U^{gH} is coarsely constant.

Consistency: Let $U, V \in \mathfrak{T}$ and $a \in G$. If $U, V \in \mathfrak{S} - \{S\}$, then consistency holds automatically. Hence suppose that $U = gH$ for some $g \in G$. If $gH \sqsubset V$, then $V = S$. In this case, $\pi_{gH}(a) = \rho_{gH}^S(\pi_S(a))$ by definition, so consistency holds. If $U \in \mathfrak{S}$ and $V = S$, then consistency follows easily from consistency in (G, \mathfrak{S}) . There is no case in which $V \sqsubset gH$.

Hence suppose $gH \triangleleft V$. Choose $b \in \rho_{gH}^S(a)$, so that b is the entrance point in $\mathfrak{C}(gH)$ of some geodesic $[a, v_{gH}]$. Then $d_U(a, \rho_U^{gH}) = d_U(a, \pi_U(gH)) \leq d_U(a, b)$. If $d_U(a, b) > E$, then ρ_U^S lies $10E$ -close to a \mathcal{CS} -geodesic from a to b , from which it is easily deduced that $d_{gH}(b, \rho_{gH}^S(\rho_U^S))$ is uniformly bounded, as required.

The bound on $d_W(\rho_W^U, \rho_W^V)$ from Definition 1.4.(4) holds automatically when $U, V, W \in \mathfrak{S}$ and holds vacuously otherwise by the definition of the nesting relation in \mathfrak{T} .

Bounded geodesic image: If $U, V \in \mathfrak{S} - \{S\}$ and $U \sqsubseteq V$, then bounded geodesic image holds because it held in (G, \mathfrak{S}) . Hence it remains to consider the case where $V = S$. First suppose that $U \in \mathfrak{S}$ and that γ is a geodesic in $\text{Pyr}(G)$ that does not pass $(E + D + 2r)$ -close to ρ_S^U . Let α be a push-off of γ and let α' be a \mathcal{CS} -geodesic at Hausdorff distance $\leq D$ from α , provided by Lemma 6.13. Then α' cannot pass through the E -neighborhood of ρ_S^U , for otherwise α would pass through the $(E + D)$ -neighborhood of ρ_S^U , whence γ would pass through the $(E + D + 2r)$ -neighborhood of ρ_S^U in $\text{Pyr}(G)$. Hence there exists $E' = E'(D, E)$ so that $\rho_V^S(\alpha')$ has diameter at most E' , by bounded geodesic image in (G, \mathfrak{S}) (here, we mean $\rho_U^S: \mathcal{CS} \rightarrow \mathcal{CU}$). Each point of γ maps by π_U to a point at distance at most C from a point of $\pi_U(\alpha')$, by Lemma 6.16, and we are done since $\rho_U^S(\mathfrak{C}(gH)) \subseteq \pi_U(gH)$ for each gH .

Next suppose that $U = gH$ for some $g \in G$. Let $x, y \in \text{Pyr}(G)$. A thin triangle argument shows that if each geodesic $[x, y]$ is sufficiently far from gH in $\text{Pyr}(G)$, then the geodesics $[x, v_{gH}], [y, v_{gH}]$ enter $\mathfrak{C}(gH)$ at points $a_x, a_y \in gH$ with $d_\Delta(a_x, a_y) \leq 100\delta$. Our choice of r ensures that $d_{gH}(a_x, a_y)$ is uniformly bounded, since points in gH at distance $< 2r$ in $\mathfrak{C}(gH)$ are at uniformly bounded distance (depending on r) in gH .

Large links: Let $a, b \in G$ and let $N = \lfloor d_\Delta(a, b) \rfloor$. We will produce uniform constants K, λ' and $T_1, \dots, T_m \in \mathfrak{S}$ and g_1H, \dots, g_nH so that $m + n \leq \lambda'N + \lambda'$ and $d_U(a, b) \leq K$ unless $U = g_iH$ or $U \sqsubseteq T_j$ for some i, j .

Fix a \mathcal{CS} -geodesic γ from a to b . By Lemma 6.13 and Lemma 6.15, for each sufficiently large $K_0 > 2r$ there exists K_1 so that either $d_{gH}(a, b) \leq K_0$ or γ has a maximal subpath γ_g lying in the K_1 -neighborhood of gH (in \mathcal{CS}), and the endpoints a_g, b_g of γ_g satisfy $d_{gH}(a, a_g) \leq K_1, d_{gH}(b, b_g) \leq K_1$. Let $G(a, b)$ be the set of gH with $d_{gH}(a, b) > K_0$. Observe that for all distinct $gH, g'H \in G(a, b)$, we have $\text{diam}_{\mathcal{CS}}(\gamma_g \cap \gamma_{g'}) \leq K_2$, where K_2 depends on K_1 and the geometric separation constants from Remark 6.3. We note that K_2 does not depend on K_0 and thus, by choosing K_0 large enough compared to K_2 , we can ensure that at most two elements of $G(a, b)$ simultaneously overlap. Observe that this implies that the cardinality of $G(a, b)$ is at most $2 \cdot d_\Delta(a, b)$.

Write $\gamma = \left(\prod_{i=1}^k \alpha_i \beta_i \right) \alpha_{k+1}$, where each β_i is a subpath contained in the union of subpaths γ_g , where $gH \in G(a, b)$, and $\text{Int}(\alpha_i)$ is in the complement of the union of such paths. By bounded geodesic image, Lemma 6.16, and the large link lemma in (G, \mathfrak{S}) , there exist $T_1, \dots, T_m \in \mathfrak{S}$ such that $d_U(a, b) \leq K_0$ unless $U \sqsubseteq T_i$ for some i , where $m = \lambda \lfloor \sum_i |\alpha_i| \rfloor + \lambda$. Hence any elements $U \in \mathfrak{T}$ in which $d_U(a, b) > \max\{E, K_0\}$ is nested into one of at most $3 \cdot (\lambda d_\Delta(a, b) + \lambda)$ elements of $\mathfrak{T} - \{S\}$.

Partial realization: Let $\{U_i\}$ be a set of pairwise-orthogonal elements of \mathfrak{T} and let $b_i \in \mathcal{C}_{\mathfrak{T}}U_i$ for each i . We consider two cases. First, if each $U_i \in \mathfrak{S}$, then partial realization in

(G, \mathfrak{S}) implies that there exists $g \in G$ so that $d_{U_i}(g, b_i) \leq E$ for each i and $d_V(b_i, \rho_V^{U_i}) \leq E$ when $U_i \sqsubseteq V$ or $U_i \pitchfork V$. Thus it remains to note that $d_{\mathcal{C}_{\mathfrak{X}}S}(b_i, \rho_V^{U_i})$ is uniformly bounded, since $\mathcal{C}S \hookrightarrow \text{Pyr}(G)$ is distance non-increasing.

When $\{U_i\} = \{S\}$, partial realization holds for $\mathcal{C}_{\mathfrak{X}}S$ since it held for $\mathcal{C}S$. Hence it remains to consider a coset gH and some $gh \in gH$. Obviously $gh \in G$ has the correct projection in gH . If $gH \sqsubseteq V$, then $V = S$ and $\rho_S^{gH} = v_{gH}$. Hence $d_{\mathcal{C}_{\mathfrak{X}}S}(gh, \rho_S^{gH}) \leq r$ as required. If $gH \pitchfork V$, then $\rho_V^{gH} = \pi_V(gH) \ni \pi_V(gh)$, as required.

Uniqueness: Let $\kappa \geq 0$ be given and let $\theta = \theta(\kappa)$ be the corresponding constant from (G, \mathfrak{S}) , so that, if $d_G(a, b) \geq \theta$ then $d_{\mathcal{C}V}(a, b) \geq \kappa$ for some $V \in \mathfrak{S}$. Hence we must consider only $a, b \in G$ such that $d_G(a, b) \geq \theta$ but $d_{\mathcal{C}V}(a, b) \leq \kappa$ if and only if $V \neq S$. Consider a geodesic γ in $\text{Pyr}(G)$ joining a to b and a push-off $\alpha = \alpha_0\beta_0 \cdots \alpha_n\beta_n\alpha_{n+1}$ of γ , where each β_i is a $\mathcal{C}S$ -geodesic joining two points in γ lying in some g_iH . Lemma 6.13 implies that α lies at Hausdorff distance D from a $\mathcal{C}S$ -geodesic α' joining a, b . By assumption, α' has length at least κ . Hence either $|\beta_i| \geq \kappa$ for some i (i.e., $d_{g_iH}(a, b) \geq \kappa$) or $n \geq \epsilon\kappa$ for some uniform $\epsilon \geq 0$, and thus $d_{\mathcal{C}_{\mathfrak{X}}S}(a, b) \geq \epsilon\kappa$. Thus for each $\kappa \geq 0$, we have that $d_V(a, b) \geq \kappa$ for some $V \in \mathfrak{X}$ provided $d_G(a, b) \geq \max\{\theta(\kappa), \theta(\epsilon^{-1}\kappa)\}$, i.e., uniqueness holds. \square

6.2. Proof of Theorem 6.2. In light of Lemma 6.13, we now enlarge r so that $r \geq 10^9 CDEQ\delta$, where C exceeds the constants from Lemma 6.15 and Lemma 6.16. Let F be a finite subset of $H - \{1\}$ chosen so that for any $N \triangleleft H$ which avoids F , yields $(\{N^g : g \in G\}, \{v_{gH} : g \in G\})$ is a *very rotating family* (see the proof of Proposition 6.7); our choice of r ensures that $\{v_{gH}\}$ is 200δ -separated. Let $\text{Pyr}(G) = \text{Pyr}(G)_r$, and let $\text{Pyr}(G/\hat{N}) = \text{Pyr}(G/\hat{N})_r$, where $N \triangleleft H$ avoids F .

6.2.1. Linked pairs.

Definition 6.17 (Fulcrum). Let $x, y \in \text{Pyr}(G)$ and $N \triangleleft H$. We say that the apex v_{gH} is a d -fulcrum for x, y if $d_{\Delta}(x, y) = d_{\Delta}(x, v_{gH}) + d_{\Delta}(v_{gH}, y)$ and there exists $h \in gNg^{-1}$, $x' \in [x, v_{gH}]$, $y' \in [v_{gH}, y]$ with the following properties.

- $d_{\Delta}(x', v_{gH}), d_{\Delta}(y', v_{gH}) \in [25\delta, 30\delta]$,
- $d_{\Delta}(x', hy') \leq d$.

A fulcrum is shown in Figure 2.




FIGURE 2. A fulcrum.

Our assumptions on N and r mean that we have the following “Greendlinger lemma” [DGO11, Lemma 5.10], which is formulated in our context as follows:

Lemma 6.18 (Greendlinger lemma). *Let $n \in \hat{N} - \{1\}$ and let $p \in \text{Pyr}(G)$. Then one of the following holds:*

- any $\text{Pyr}(G)$ -geodesic $[p, np]$ contains a δ -fulcrum for p, np ; or,
- there exists v_{gH} so that $n \in gNg^{-1} \triangleleft gHg^{-1}$ and $d_{\Delta}(p, v_{gH}) \leq 25\delta$.

In [DGO11], the first conclusion uses 5δ , but in fact the conclusion can be made to hold for an arbitrarily small constant, and we use δ for convenience.

Remark 6.19. For convenience, we can and shall assume that, for all $U \in \mathfrak{S} - \{S\}$ and all $x \in G$, the sets ρ_S^U and $\pi_S(x)$ consist of single points, and $\rho_S^{gU} = g\rho_S^U$ and $\pi_S(gx) = g\pi_S(x)$ for all $g \in G$. Indeed, equivariantly replace each relevant bounded set with one of its elements, and adjust the constants of Definition 1.4 and Subsection 1.2.3 uniformly if necessary.

Definition 6.20 (Linked pair). Let $U, V \in \mathfrak{S}$ (resp. $x, y \in G$, resp. $U \in \mathfrak{S}, x \in G$). Then $\{U, V\}$ (resp. $\{x, y\}$, resp. $\{U, x\}$) is *linked* if there does not exist a 10δ -fulcrum for ρ_S^U, ρ_S^V (resp. $\pi_S(x), \pi_S(y)$, resp. $\rho_S^U, \pi_S(x)$). We say the pair is *weakly linked* when there is no 5δ -fulcrum.

Lemma 6.21. *Linked pairs have the following properties:*

- (1) for all $[U], [V] \in \mathfrak{S}/\widehat{N}$, there exists a (weakly) linked pair $U \in [U], V \in [V]$, and the same holds for pairs $\widehat{N}x, \widehat{N}y$;
- (2) for any $g \in G$ the pair $\{gU, gV\}$ is (weakly) linked whenever $\{U, V\}$ is (weakly) linked; the same holds for linked pairs $\{x, y\}$;
- (3) if $x \in G$ and $n \in \widehat{N}$, then $\{x, nx\}$ are not weakly linked and, in fact, have a δ -fulcrum (and the same holds for x and xn since \widehat{N} is normal).

Proof. Choose $x, y \in \text{Pyr}(G)$ and suppose that v_{gH} is a 5δ -fulcrum for x, y (i.e., x, y are not weakly linked). Choose $x', y' \in [x, v_{gH}], [y, v_{gH}]$ and $h \in gHg^{-1}$ so that $d_\Delta(x', v_{gH}), d_\Delta(y', v_{gH}) \in [25\delta, 30\delta]$ and $d_\Delta(hy', x') \leq 5\delta$. Then

$$d_\Delta(x, hy) \leq d_\Delta(x, x') + d_\Delta(y, y') + 5\delta \leq d_\Delta(x, y) - 45\delta,$$

so, by replacing x, y with x, hy , we obtain a closer pair of representatives of $\widehat{N}x, \widehat{N}y$. This proves assertion (1) for weakly linked pairs. Repeating exactly the same argument with 10δ replacing 5δ establishes the assertion for linked pairs.

For all $x, y \in \text{Pyr}(G)$, cosets $g'H$, and $g \in G$, and $d \geq 0$, observe that $v_{g'H}$ is a d -fulcrum for x, y if and only if $v_{gg'H}$ is a d -fulcrum for gx, gy , which proves assertion (2). Assertion (3) follows from Lemma 6.18. \square

6.2.2. Proof of Theorem 6.2. Throughout this section, we say that $N \triangleleft H$ is *sufficiently deep* if $N \cap F = \emptyset$, where $F \subset H - \{1\}$ is the finite subset whose exclusion from N implies that $(\{gNg^{-1}\}, \{v_{gH}\})$ is a 200δ -separated very rotating family. We now define the hierarchical space structure on G/\widehat{N} . By choosing N sufficiently deep, we can ensure that, if $g, g' \in G$ differ by a generator in a fixed finite generating set, then g, g' are linked.

Construction 6.22. The index set and associated hyperbolic spaces are defined by the following, where (G, \mathfrak{S}) is the original HHG structure and the modified HHG structure provided by Proposition 6.14 is denoted (G, \mathfrak{T}) :

- (1) the index set \mathfrak{S}_N is \mathfrak{T}/\widehat{N} ;
- (2) for $\mathbf{S} = \{S\}$ (the \widehat{N} -orbit of S), let $\mathcal{CS} = \text{Cay}(G/\widehat{N}, \mathcal{T}/\widehat{N} \cup H/N)$ — note that this is quasi-isometric to $\text{Pyr}(G/\widehat{N})$;
- (3) $\mathbf{U} \in \mathfrak{S}/\widehat{N}$, let $\mathcal{CU} = (\bigsqcup_{U \in \mathbf{U}} \mathcal{CU})/\widehat{N}$ — note that this is isomorphic to \mathcal{CU} for some (hence any) $U \in \mathbf{U}$;
- (4) for $\mathbf{U} = \widehat{N}gH$, let $\mathcal{CU} = \left(\bigsqcup_{ngH \in \mathbf{U}} \text{Cay}(H, \mathcal{T} \cap H) \right) / \widehat{N}$ — note that this is isometric to a Cayley graph of H/N ;
- (5) for each $\mathbf{U}, \mathbf{V} \in \mathfrak{S}_N$, let $\mathbf{U} \sqsubseteq \mathbf{V}$ (resp. $\mathbf{U} \perp \mathbf{V}$) if there exists a linked pair $\{U, V\} \subseteq \mathfrak{S}$ with $U \in \mathbf{U}, V \in \mathbf{V}$ so that $U \sqsubseteq V$ (resp. $U \perp V$). If $\mathbf{U} = \widehat{N}gH$, then we let $\mathbf{U} \sqsubseteq \mathbf{S}$. If neither $\mathbf{U} \sqsubseteq \mathbf{V}, \mathbf{V} \sqsubseteq \mathbf{U}$ nor $\mathbf{U} \perp \mathbf{V}$ holds, then we let $\mathbf{U} \not\sqsubseteq \mathbf{V}$.

The projections are defined by taking all linked pair representatives:

- (6) $\pi_{\mathbf{U}}(g\hat{N}) = (\bigcup \pi_{U'}(g')) / \hat{N}$, where the union is taken over all linked pairs $\{U', g'\}$ with $U' \in \mathbf{U}$, $g' \in g\hat{N}$;
- (7) similarly, for $\mathbf{U} \sqsubset \mathbf{V}$ or $\mathbf{U} \pitchfork \mathbf{V}$, $\rho_{\mathbf{V}}^{\mathbf{U}}$ is defined as $(\bigcup \rho_{V'}^{U'}) / \hat{N}$, where the union is taken over all linked pairs $\{U', V'\}$ with $U' \in \mathbf{U}$, $V' \in \mathbf{V}$;
- (8) finally, for $\mathbf{V} \sqsubset \mathbf{U}$ and $\hat{N}x \in \mathcal{C}\mathbf{U}$, let $\rho_{\mathbf{V}}^{\mathbf{U}}(\hat{N}x) = (\bigcup \rho_{V'}^{U'}(x')) / \hat{N}$, where the union is taken over all linked pairs $\{U', V'\}$ with $U' \in \mathbf{U}$, $V' \in \mathbf{V}$ and $x' \in \hat{N}x \cap \mathcal{C}U'$.

Before proceeding with the proof of Theorem 6.2, we need several lemmas:

Lemma 6.23. *If N is sufficiently deep, then for any $\mathbf{U}, \mathbf{V} \in \mathfrak{S}/\hat{N} - \{\mathbf{S}_N\}$ and any $U \in \mathbf{U}$ there exists at most one $V \in \mathbf{V}$ with $d_{\mathcal{C}\mathbf{S}}(\rho_S^U, \rho_S^V) \leq 10E$, and hence in particular at most one such V with $U \sqsubseteq V$, $V \sqsubseteq U$ or $U \perp V$.*

Proof. It follows from Lemma 6.18 that, for each $n \in N - \{1\}$ and $x \in \mathcal{C}\mathbf{S}$, we have $d_{\Delta}(x, nx) \geq 2r$, and hence $d_{\mathcal{C}\mathbf{S}}(x, nx) \geq 2r > 10E$. The last assertion follows from Definition 1.4.4 when $U \sqsubseteq V$ or $V \sqsubseteq U$, and from Lemma 1.8 when $U \perp V$; in both cases ρ_S^U and ρ_S^V are E -close. \square

Lemma 6.24. *If N is sufficiently deep, then the following holds for each $x, y \in G$: if $\{U, \pi_S(x)\}$ and $\{U, \pi_S(y)\}$ are linked (resp. weakly linked) then either:*

- (1) $\{\pi_S(x), \pi_S(y)\}$ is weakly linked (resp. has no δ -fulcrum), or
- (2) there exists v_{gH} with $d_{\Delta}([\pi_S(x), \rho_S^U], v_{gH}) \leq 40\delta$ and $d_{\Delta}([\pi_S(y), \rho_S^U], v_{gH}) \leq 40\delta$.

The analogous statement holds when replacing x and/or y with an element of $\mathfrak{S} - \{S\}$.

Proof. Let U, x, y be as in the statement and suppose that $\{x, y\}$ is not weakly linked. By definition, there exists a 5δ -fulcrum v_{gH} for $\pi_S(x), \pi_S(y)$. Consider a geodesic triangle in the δ -hyperbolic space $\text{Pyr}(G)$ with vertices $\pi_S(x), \pi_S(y)$ and ρ_S^U , with $v_{gH} \in [\pi_S(x), \pi_S(y)]$. Choose $x' \in [x, v_{gH}]$, $y' \in [v_{gH}, y]$ so that $d_{\Delta}(x', v_{gH}), d_{\Delta}(y', v_{gH}) \in [25\delta, 30\delta]$ and so that $d_{\Delta}(x', hy') \leq 5\delta$ for some $h \in gHg^{-1}$.




FIGURE 3. Finding an illegal fulcrum in the proof of Lemma 6.24.

If $d_{\Delta}(x', [\rho_S^U, \pi_S(x)]), d_{\Delta}(y', [\rho_S^U, \pi_S(x)]) > \delta$, then the very rotating condition [DGO11, Lemma 5.5] implies that v_{gH} is contained in $[\rho_S^U, \pi_S(y)]$ and that v_{gH} is a 10δ -fulcrum for $\pi_S(y), \rho_S^U$ (witnessed by the same element $h \in gHg^{-1}$), contradicting that U, y are linked. (See Figure 3.) Since the same argument works for $[\rho_S^U, \pi_S(y)]$, we have

$$d_{\Delta}([\rho_S^U, \pi_S(x)], \{x', y'\}) \leq \delta, \quad d_{\Delta}([\rho_S^U, \pi_S(y)], \{x', y'\}) \leq \delta,$$

and $d_{\Delta}(v_{gH}, x'), d_{\Delta}(v_{gH}, y') \leq 30\delta$, so the claim follows. \square

Lemma 6.25. *There exists $K \geq 0$ so that the following holds: if $x, g \in G$ and $U \in \mathfrak{S}$ satisfy $d_{\Delta}([\pi_S(x), \rho_S^U], v_{gH}) \leq 40\delta$ then $\text{diam}_{\mathcal{C}U}(\pi_U(x) \cup gH) \leq K$.*

Proof. Observe that $[\pi_S(x), \rho_S^U]$ must pass through $\mathfrak{C}(gH)$, since $d_\Delta([\pi_S(x), \rho_S^U], v_{gH}) \leq 40\delta$ and $r \geq 10^9\delta$. Let a be the entry point of $[x, \rho_S^U]$ in $\mathfrak{C}(gH)$ and let $[x, a]$ be the sub-geodesic joining x to a . Let $[a, b]$ be the sub-geodesic of $[\pi_S(x), \rho_S^U]$ that joins the entry point a of $[\pi_S(x), \rho_S^U]$ in $\mathfrak{C}(gH)$ to the entry point b of $[\rho_S^U, x]$ in $\mathfrak{C}(gH)$. Since $d_\Delta(a, b) \geq 2r - 80\delta > 1000CDE\delta$, we have $d_{CS}(a, b) > 1000CDE\delta$. Combined with Lemma 6.13, this shows that $d_{CS}(\rho_S^U, c) > 100E$ for any c on a CS -geodesic from x to a . Hence $d_{CS}([x, a], \rho_S^U) > E$, so bounded geodesic image in (G, \mathfrak{S}) implies that $d_U(x, a) \leq E$, whence $\text{diam}_U(\pi_U(x) \cup \pi_U(gH)) \leq E + C$, by Lemma 6.16. \square

Combining Lemma 6.24 and Lemma 6.25, and increasing C if necessary, yields:

Corollary 6.26. *There exists $C \geq 0$ so that the following holds for each $x, y \in G$ provided N is sufficiently deep. If $\{U, \pi_S(x)\}$ and $\{U, \pi_S(y)\}$ are linked (resp. weakly linked) but $\{\pi_S(x), \pi_S(y)\}$ is not weakly linked (resp. has a δ -fulcrum), then $d_{CU}(x, y) \leq C$. The same holds with x and/or y replaced with elements of $\mathfrak{S} - \{S\}$.*

Lemma 6.27. *Let $\{\mathbf{U}_i\}_{i=1}^k$ be a totally orthogonal set with $\mathbf{U}_i \in \mathfrak{S}/\widehat{N}$ for all i . Then there exist representatives $U_i \in \mathbf{U}_i$ so that for all distinct i, j , we have $U_i \perp U_j$ and $\{U_i, U_j\}$ is a linked pair.*

Proof. This follows by induction on k , using Lemma 6.28. Indeed, when $k = 1$, there is nothing to prove. Suppose that we can choose the required U_i for $1 \leq i \leq k - 1$. For each $i \leq k - 1$, choose $U_k^i \in \mathbf{U}_k$ so that $\{U_i, U_k^i\}$ is a linked pair and $U_i \perp U_k^i$. Then for each i, j , we have $d_\Delta(\rho_S^{U_k^i}, \rho_S^{U_k^j}) \leq 30E$, so, by Lemma 6.28, the U_k^i all coincide, and we are done. \square

Lemma 6.28. *Let $U \in \mathfrak{S}$ and let $n \in \widehat{N}$. Then either $nU = U$ or $d_\Delta(\rho_S^U, \rho_S^{nU}) > 100E$.*

Proof. If $U \neq nU$ and $d_\Delta(\rho_S^U, \rho_S^{nU}) \leq 100E$, then each geodesic $[\rho_S^U, \rho_S^{nU}]$ in $\text{Pyr}(G)$ fails to pass through any apex, since $r > 10^9E$. In particular, there is no 5δ -fulcrum for $\rho_S^U, \rho_S^{nU} = n\rho_S^U$. Thus, by Lemma 6.18, there exists v_{gH} so that $n \in gNg^{-1}$ and $d_\Delta(\rho_S^U, v_{gH}) \leq 25\delta$. But this is impossible, since $\rho_S^U \in CS$ lies at distance at least $r > 10^9\delta$ from any apex. \square

Lemma 6.29. *There exists C' so that the following holds. Let $\{g, g'\}$ be d_Δ -minimal representatives of $g\widehat{N}, g'\widehat{N}$ and let $U \in \mathfrak{T}$ satisfy $d_U(g, g') > C'$. Then $\{U, g\}$ and $\{U, g'\}$ are linked.*

Proof. Consider a geodesic triangle in $\text{Pyr}(G)$ formed by g, g', ρ_S^U . Suppose $\{U, g\}$ is not linked, so $[\rho_S^U, g']$ contains a 10δ -fulcrum $v = v_{g''H}$ for $\{U, g'\}$. If $[g, \rho_S^U]$ passes 40δ -close to v , then Lemma 6.25 shows that $\pi_U(g), \pi_U(g')$ coarsely coincide with $\rho_U^{g''H}$. Otherwise, $[\rho_S^U, g']$ contains a length- 60δ subpath, centered at v and contained in $\mathcal{N}_\delta([g, g'])$. Using the notation of Definition 6.17, let $h \in N^{g''}, x', y' \in [\rho_S^U, g]$ witness the fact that v is a 10δ -fulcrum for ρ_S^U, g' , with y' between v and g' . Choose $x'', y'' \in [g, g']$ with $d_\Delta(x', x''), d_\Delta(y', y'') \leq \delta$. Then $d_\Delta(g, hg') \leq d_\Delta(g, x'') + d_\Delta(y'', g') + 12\delta < d_\Delta(g, g')$, contradicting our choice of g, g' . \square

We are now ready to complete the proof of Theorem 6.2:

Proof of Theorem 6.2. The claimed hierarchical space structure $(G/\widehat{N}, \mathfrak{S}_N)$ is described in Construction 6.22. Observe that each CU is uniformly hyperbolic by definition when $\mathbf{U} = \widehat{N}U$ for some $U \in \mathfrak{S} - \{S\}$. Moreover, CS is hyperbolic by Proposition 6.7. If $\mathbf{U} \in \mathfrak{S}_N$ arose from a coset of H/N , then \mathbf{U} is necessarily \sqsubseteq -minimal. Hence, if $(G/\widehat{N}, \mathfrak{S}_N)$ is a hierarchical space structure, then it is a relatively hierarchically hyperbolic space structure. Moreover, G/\widehat{N} acts on \mathfrak{S}_N and, for each $g \in G/\widehat{N}$ and $\mathbf{U} \in \mathfrak{S}_N$, it is easily seen that there is an induced

isometry $\mathcal{C}\mathbf{U} \rightarrow \mathcal{C}g\mathbf{U}$ so that the required diagrams from Section 1.2.3 coarsely commute. Hence it suffices to show that $(G/\widehat{N}, \mathfrak{S}_N)$ is a hierarchical space.

We observe that if $\{U, x\}$ is linked and $\{U, nx\}$ is weakly linked, then $d_U(x, nx) \leq C$, for C as in Corollary 6.26. In fact, $\{x, nx\}$ has a δ -fuchrum by Lemma 6.18.

Verifying Definition 1.4.(1): To finish proving that $(G/\widehat{N}, \mathfrak{S}_N)$ satisfies the projections axiom, we must check that each $\pi_{\mathbf{U}}$ sends points to uniformly bounded sets and is uniformly coarsely Lipschitz. Let $\mathbf{U} \in \mathfrak{S}_N$, then for each $g \in G$ we have that $\pi_{\mathbf{U}}(g\widehat{N})$ is uniformly bounded by Corollary 6.26 and Lemma 6.21.(3).

We now show that $\pi_{\mathbf{U}}$ is coarsely Lipschitz. Let $g\widehat{N}, g'\widehat{N} \in G/\widehat{N}$. It suffices to consider the case where $d_{G/\widehat{N}}(g\widehat{N}, g'\widehat{N}) \leq 1$. Choose representatives g, g' that differ by a single generator of G . Let $\mathbf{U} \in \mathfrak{T}$. Choose $U \in \mathbf{U}$ so that $\{U, g\}$ is a linked pair, so that since N was chosen sufficiently deep, $\{U, g'\}$ is weakly linked. Hence, the observation above shows $d_U(g, g') \leq C$, as required.

Verifying Definition 1.4.(2),(3),(5),(7): The nesting, orthogonality, finite complexity and bounded geodesic image axioms easily follow from Lemma 6.23.

Verifying Definition 1.4.(4): We now prove consistency. Let $\mathbf{U} \triangleleft \mathbf{V}$ and $g\widehat{N} \in G/\widehat{N}$. Also, suppose $U \in \mathbf{U}, V \in \mathbf{V}$ and g have the property that $\{U, V\}$ and $\{U, g\}$ are linked; this is justified by Lemma 6.21. If $\{V, g\}$ is weakly linked, then, using the observation above, consistency for $\mathbf{U}, \mathbf{V}, g\widehat{N}$ follows from consistency for U, V, g (Proposition 6.14). If not, by Corollary 6.26 we have $d_U(\rho_U^V, g) \leq C$, and hence the consistency inequality holds.

Now suppose that $\mathbf{U} \not\sqsubseteq \mathbf{V}$ and $g\widehat{N} \in G/\widehat{N}$. Also, suppose $U \in \mathbf{U}, V \in \mathbf{V}$ and g have the property that $\{U, V\}$ and $\{U, g\}$ are linked and $U \not\sqsubseteq V$. If $\{V, g\}$ is a weakly linked pair, then consistency follows from consistency for U, V, g . Otherwise, apply Corollary 6.26 as above.

If $\mathbf{U} \sqsubseteq \mathbf{V}$ and \mathbf{W} satisfies either $\mathbf{V} \not\sqsubseteq \mathbf{W}$ or $\mathbf{V} \triangleleft \mathbf{W}$ and $\mathbf{W} \not\perp \mathbf{U}$, then $d_{\mathbf{W}}(\rho_{\mathbf{W}}^{\mathbf{U}}, \rho_{\mathbf{W}}^{\mathbf{V}})$ is uniformly bounded by a similar argument. This completes the proof of consistency.

Verifying Definition 1.4.(6): Let $g\widehat{N}, g'\widehat{N} \in G/\widehat{N}$ and let $\mathbf{W} \in \mathfrak{S}_N$. We divide into three cases according to whether $\mathcal{C}\mathbf{W} = g''H/N$ for some g'' , or $\mathcal{C}\mathbf{W} = CW$ for some $W \in \mathbf{W}$, or $\mathcal{C}\mathbf{W} = \text{Pyr}(G/\widehat{N})$. In the first case, nothing is properly nested into \mathbf{W} and we are done.

Consider the second case, and choose $g, g' \in g\widehat{N}, g'\widehat{N}$ and $W, W' \in \mathbf{W}$ so that $\{W, g\}$ and $\{W', g'\}$ are linked pairs. By translating, we may assume that $W = W'$, and $d_{\mathbf{W}}(g\widehat{N}, g'\widehat{N}) = d_W(g, g')$ by definition. Hence the large link lemma in (G, \mathfrak{T}) provides $\mathbf{T}_1, \dots, \mathbf{T}_k \sqsubseteq \mathbf{W}$, with \mathbf{T}_i represented by some $T_i \in \mathfrak{S}$ so that if $U \not\sqsubseteq W$, then $d_U(g, g') > E$ only if $U \sqsubseteq T_i$ for some i . Suppose that $d_{\mathbf{U}}(g\widehat{N}, g'\widehat{N}) > E + 2C$ for some $\mathbf{U} \not\sqsubseteq \mathbf{W}$. Then there exist $g_1, g'_1 \in g\widehat{N}, g'\widehat{N}$ so that g_1, g'_1 are both linked to some $U \in \mathbf{U}$ with $U \not\sqsubseteq W$. Lemma 6.21.(3) and Corollary 6.26 imply that $\pi_U(g), \pi_U(g_1)$ C -coarsely coincide, and the same is true for $\pi_U(g'), \pi_U(g'_1)$, so U must be nested in some T_i . Since $\rho_{\mathfrak{S}}^U$ and $\rho_{\mathfrak{S}}^{T_i}$ coarsely coincide, it follows that U, T_i is a linked pair and thus $\mathbf{U} \sqsubseteq \mathbf{T}_i$.

Consider the third case. Let $g, g' \in g\widehat{N}, g'\widehat{N}$ be minimal-distance (in d_{Δ}) representatives. Observe that $\{g, g'\}$ is a linked pair and that $d_{\Delta}(g, g') = d_{\Delta_N}(g\widehat{N}, g'\widehat{N})$, where d_{Δ_N} is the metric on $\text{Pyr}(G/\widehat{N})$. The claim follows from the large link lemma in (G, \mathfrak{T}) as above.

Verifying Definition 1.4.(8): Let $\{\mathbf{U}_i\}_{i=1}^k$ be a totally orthogonal subset of \mathfrak{T}/\widehat{N} . If $\mathbf{U}_i \notin \mathfrak{S}/\widehat{N}$ for some i , then $k = 1$ and partial realization obviously holds. Hence suppose that $\mathbf{U}_i \in \mathfrak{S}/\widehat{N}$ for all i . Then, by Lemma 6.27, for each $i \leq k$, there exists $U_i \in \mathbf{U}_i$ so that for all distinct i, j , we have $U_i \perp U_j$ and $\{U_i, U_j\}$ is a linked pair. The claim now follows from partial realization in (G, \mathfrak{T}) .

Verifying Definition 1.4.(9): Let $g\widehat{N}, g'\widehat{N} \in G/\widehat{N}$ and let $\kappa \geq 0$. Suppose that for all $U \in \mathfrak{T}$, we have $d_U(g\widehat{N}, g'\widehat{N}) \leq \kappa$. Let $g, g' \in g\widehat{N}, g'\widehat{N}$ be minimal-distance (in d_Δ) representatives.

We now show that $\pi_U(g), \pi_U(g')$ are $(\kappa + C')$ -close in every $U \in \mathfrak{S}$, for C' as in Lemma 6.29. By uniqueness in (G, \mathfrak{T}) , it follows that $d_G(g, g') \leq \theta(\kappa + C')$, which implies the required bound on $d_{G/\widehat{N}}(g\widehat{N}, g'\widehat{N})$.

Since we chose g, g' at minimal distance, we have $d_\Delta(g, g') \leq \kappa$. Suppose that there exists $U \in \mathfrak{T} - \{S\}$ with $d_U(g, g') > C'$. Then by Lemma 6.29, U is linked to both g and g' , and hence $d_U(g, g') = d_U(g\widehat{N}, g'\widehat{N}) \leq \kappa$, as required. \square

REFERENCES

- [APMS15] Yago Antolin Pichel, Ashot Minasyan, and Alessandro Sisto. Commensurating endomorphisms of acylindrically hyperbolic groups and applications. *Groups Geometry and Dynamics*, pages 1–46, 2015. To appear.
- [BB15] Mladen Bestvina and Ken Bromberg. On the asymptotic dimension of the curve complex. *arXiv e-print 1508.04832*, 2015.
- [BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math. Inst. Hautes Études Sci.*, 122:1–64, 2015.
- [BD01] Gregory Bell and Alexander Dranishnikov. On asymptotic dimension of groups. *Algebraic & Geometric Topology*, 1(1):57–71, 2001.
- [BD06] Gregory Bell and Alexander Dranishnikov. A Hurewicz-type theorem for asymptotic dimension and applications to geometric group theory. *Transactions of the American Mathematical Society*, 358(11):4749–4764, 2006.
- [BD08] Gregory Bell and Alexander Dranishnikov. Asymptotic dimension. *Topology and its Applications*, 155(12):1265–1296, 2008.
- [BD11] Gregory Bell and Alexander Dranishnikov. Asymptotic dimension in Będlewo. *Topology Proc.*, 38:209–236, 2011.
- [BF07] Gregory Bell and Koji Fujiwara. The asymptotic dimension of a curve graph is finite. *Journal of the London Mathematical Society*, 2007.
- [BHS14] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces I: curve complexes for cubical groups. *arXiv:1412.2171*, 2014.
- [BHS15] Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: combination theorems and the distance formula. *arXiv eprint 1509.00632v2*, pages 1–53, 2015.
- [BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. *Geom. Topol.*, 16(2):781–888, 2012.
- [BM00] Marc Burger and Shahar Mozes. Lattices in product of trees. *Inst. Hautes Études Sci. Publ. Math.*, 92(1):151–194, 2000.
- [Bow08] Brian H. Bowditch. Tight geodesics in the curve complex. *Inventiones mathematicae*, 171(2):281–300, 2008.
- [CdH14] Yves Cornuier and Pierre de la Harpe. Metric geometry of locally compact groups. *arXiv preprint arXiv:1403.3796*, 2014.
- [DGO11] François Dahmani, Vincent Guirardel, and Denis Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *arXiv preprint arXiv:1111.7048*, 2011.
- [DHS16] Matthew G Durham, Mark F Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *arXiv preprint arXiv:1604.01061*, 2016.
- [DS05] Cornelia Druțu and Mark Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005.
- [GM08] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel Journal of Mathematics*, 168(1):317–429, 2008.
- [Gro96] Mikhail Gromov. Geometric group theory, vol. 2: Asymptotic invariants of infinite groups. *Bull. Amer. Math. Soc.*, 33:0273–0979, 1996.
- [Gui14] Vincent Guirardel. Geometric small cancellation. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 55–90. Amer. Math. Soc., Providence, RI, 2014.
- [Hag13] Mark F. Hagen. The simplicial boundary of a CAT(0) cube complex. *Alg. Geom. Topol.*, 13:1299–1367, 2013.

- [Hag14] Mark F. Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.*, 7(2):385–418, 2014.
- [HS16] Mark F Hagen and Tim Susse. Hierarchical hyperbolicity of all cubical groups. *arXiv preprint arXiv:1609.01313*, 2016.
- [Hul13] M Hull. Small cancellation in acylindrically hyperbolic groups. *arXiv preprint arXiv:1308.4345*, 2013.
- [MM99] Howard A Masur and Yair N Minsky. Geometry of the complex of curves I: Hyperbolicity. *Inventiones mathematicae*, 138(1):103–149, 1999.
- [MM00] Howard A Masur and Yair N Minsky. Geometry of the complex of curves II: Hierarchical structure. *Geometric and Functional Analysis*, 10(4):902–974, 2000.
- [Osi05] Denis Osin. Asymptotic dimension of relatively hyperbolic groups. *International Mathematics Research Notices*, 2005(35):2143–2161, 2005.
- [Osi07] Denis V Osin. Peripheral fillings of relatively hyperbolic groups. *Inventiones mathematicae*, 167(2):295–326, 2007.
- [Sis12] Alessandro Sisto. On metric relative hyperbolicity. *arXiv preprint arXiv:1210.8081*, 2012.
- [Smi10] A. Smirnov. The linearly controlled asymptotic dimension of the fundamental group of a graph manifold. *Algebra i Analiz*, 22(2):185–203, 2010.
- [Web15] Richard C. H. Webb. Combinatorics of tight geodesics and stable lengths. *Trans. Amer. Math. Soc.*, 367(10):7323–7342, 2015.
- [Wis] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy. 205 pp. Preprint 2011.
- [Wis07] Daniel T Wise. Complete square complexes. *Commentarii Mathematici Helvetici*, 82(4):683–724, 2007.
- [Wri12] Nick Wright. Finite asymptotic dimension for CAT(0) cube complexes. *Geometry & Topology*, 16(1):527–554, 2012.
- [Yu98] Guoliang Yu. The Novikov conjecture for groups with finite asymptotic dimension. *Annals of Mathematics*, pages 325–355, 1998.

LEHMAN COLLEGE AND THE GRADUATE CENTER, CUNY, NEW YORK, NEW YORK, USA
E-mail address: jason.behrstock@lehman.cuny.edu

DEPT. OF PURE MATHS AND MATH. STAT., UNIVERSITY OF CAMBRIDGE, CAMBRIDGE, UK
E-mail address: markfhagen@gmail.com

ETH, ZÜRICH, SWITZERLAND
E-mail address: sisto@math.ethz.ch