LARGEST ACYLINDRICAL ACTIONS AND STABILITY IN HIERARCHICALLY HYPERBOLIC GROUPS

CAROLYN ABBOTT, JASON BEHRSTOCK, AND MATTHEW GENTRY DURHAM

With an appendix by DANIEL BERLYNE and JACOB RUSSELL

Abstract. We consider two manifestations of non-positive curvature: acylindrical actions (on hyperbolic spaces) and quasigeodesic stability. We study these properties for the class of hierarchically hyperbolic groups, which is a general framework for simultaneously studying many important families of groups, including mapping class groups, right-angled Coxeter groups, most 3–manifold groups, right-angled Artin groups, and many others.

A group that admits an acylindrical action on a hyperbolic space may admit many such actions on different hyperbolic spaces. It is natural to try to develop an understanding of all such actions and to search for a "best" one. The set of all cobounded acylindrical actions on hyperbolic spaces admits a natural poset structure, and in this paper we prove that all hierarchically hyperbolic groups admit a unique action which is the largest in this poset. The action we construct is also universal in the sense that every element which acts loxodromically in some acylindrical action on a hyperbolic space does so in this one. Special cases of this result are themselves new and interesting. For instance, this is the first proof that right-angled Coxeter groups admit universal acylindrical actions.

The notion of quasigeodesic stability of subgroups provides a natural analogue of quasiconvexity which can be considered outside the context of hyperbolic groups. In this paper, we provide a complete classification of stable subgroups of hierarchically hyperbolic groups, generalizing and extending results that are known in the context of mapping class groups and right-angled Artin groups. Along the way, we provide a characterization of contracting quasigeodesics; interestingly, in this generality the proof is much simpler than in the special cases where it was already known.

In the appendix, it is verified that any space satisfying the *a priori* weaker property of being an "almost hierarchically hyperbolic space" is actually a hierarchically hyperbolic space. The results of the appendix are used to streamline the proofs in the main text.

Contents

1. Introduction 1
2. Background 5
3. Altering the hierarchically hyperbolic structure 11
4. Characterization of contracting geodesics 15
5. Universal and largest acylindrical actions 16
6. Characterizing stability 20
7. Clean containers 24
Appendix A. Almost HHSs are HHSs. By Daniel Berlyne and Jacob Russell 26
References 33

1. Introduction

Hierarchically hyperbolic groups were recently introduced by Behrstock, Hagen, and Sisto [BHS17b] to provide a uniform framework in which to study many important families of groups, including mapping class groups of finite type surfaces, right-angled Coxeter groups, most 3–manifold groups, right-angled Artin groups and many others. A *hierarchically hyperbolic space* consists of: a quasigeodesic space, $\mathcal{X}$; a set of *domains*, $\mathfrak{S}$, which index a collection of δ–hyperbolic spaces to which $\mathcal{X}$ projects; and, some additional information about these projections, including, for instance, a partial order on the domains and a unique largest element in that order, which we denote by S (i.e., S is comparable to and larger than every other domain in $\mathfrak{S}$).

Largest acylindrical actions. The study of acylindrical actions on hyperbolic spaces, as initiated in its current form by Osin [Osi16] building on earlier work of Sela [Sel97] and Bowditch [Bow08], has proven to be a powerful tool for studying groups with some aspects of non-positive curvature. As established in [BHS17b], non-virtually cyclic hierarchically hyperbolic groups admit non-elementary acylindrical actions when the δ–hyperbolic space associated to the maximal element in $\mathfrak{S}$ has infinite diameter, a property which holds in all the above examples except for those that are direct products.

Any given group with an acylindrical action may actually admit many acylindrical actions on many different spaces. A natural question is to try and find a "best" acylindrical action. There are different ways that one might try to optimize the acylindrical action. For instance, the notion of a *universal acylindrical action*, for a given group G, is an acylindrical action on a hyperbolic space X such that every element of G which acts loxodromically in some acylindrical action on some hyperbolic space, must act loxodromically in its action on X. As established by Abbott, there exist finitely generated groups which admit acylindrical actions, but no universal acylindrical action [Abb16]; we also note that universal actions need not be unique [ABO19].

In [ABO19], Abbott, Balasubramanya, and Osin introduce a partial order on cobounded acylindrical actions which, in a certain sense, encodes how much information the action provides about the group. When there exists an element in this partial ordering which is comparable to and larger than all other elements it is called a *largest* action. By construction, any largest action is necessarily a universal acylindrical action and unique.

In this paper we construct a largest action for every hierarchically hyperbolic group. Special cases of this theorem recover some recent results of [ABO19], as well as a number of new cases. For instance, in the case of right-angled Coxeter groups (and more generally for special cubulated groups), even the existence of a universal acylindrical action was unknown. Further, outside of the relatively hyperbolic setting, our result provides a single construction that simultaneously covers these new cases as well as all previously known largest and universal acylindrical actions of finitely presented groups. The following summarizes the main results of Section 5 (where there are also further details on the background and comparison with known results).

Theorem A (HHG have actions that are largest and universal)**.** *Every hierarchically hyperbolic group admits a largest acylindrical action. In particular, the following admit acylindrical actions which are largest and universal:*

(1) Hyperbolic groups.
(2) Mapping class groups of orientable surfaces of finite type.
(3) Fundamental groups of compact three-manifolds with no Nil or Sol in their prime decomposition.
(4) Groups that act properly and cocompactly on a special CAT(0) cube complex, and more generally any cubical group which admits a factor system. This includes right angled-Artin groups, right-angled Coxeter groups, and many other examples as in [HS16].

We use this construction of a largest action to characterize stable subgroups (Theorem B) and contracting elements (Corollary 5.5) of hierarchically hyperbolic groups, and to describe random subgroups of hierarchically hyperbolic groups (Theorem E).

Stability in hierarchically hyperbolic groups. One of the key features of a Gromov hyperbolic space is that every geodesic is uniformly *Morse*, a property also known as *(quasi-geodesic) stability*; that is, any uniform quasigeodesic beginning and ending on a geodesic must lie uniformly close to it. In fact, any geodesic metric space in which each geodesic is uniformly Morse is hyperbolic.

In the context of geodesic metric spaces, the presence of Morse geodesics has important structural consequences for the space; for instance, any asymptotic cone of such a space has global cut points [DMS10]. Moreover, quasigeodesic stability in groups is quite prevalent, since any finitely generated acylindrically hyperbolic group contains Morse geodesics [Osi16, Sis16].

There has been much interest in developing alternative characterizations [DMS10, CS15, ACGH17, ADT17] and understanding this phenomenon in various important contexts [Min96, Beh06, DMS10, DT15, ADT17]. This includes the theory of Morse boundaries, which encode all Morse geodesics of a group [CS15, Cor17, CH17, CD19, CM19]. In [DT15], Durham and Taylor generalized the notion of stability to subspaces and subgroups.

In this paper, we obtain a complete characterization of stability in hierarchically hyperbolic groups.

Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. We say that a subset $\mathcal{Y} \subset \mathcal{X}$ has *D–bounded projections* when $\mathsf{diam}_{\mathcal{C}U}(\pi_U(\mathcal{Y})) < D$ for all non-maximal $U \in \mathfrak{S}$; when the constant does not matter, we simply say the subset has *uniformly bounded projections.*

Theorem B (Equivalent conditions for subgroup stability)**.** *Any hierarchically hyperbolic group G admits a hierarchically hyperbolic group structure $(G, \mathfrak{S})$ such that for any finitely generated $H < G$, the following are equivalent:*

(1) H is stable in G;
(2) H is undistorted in G and has uniformly bounded projections;
(3) Any orbit map $H \to \mathcal{C}S$ is a quasi-isometric embedding, where S is the $\sqsubseteq$–maximal element in $\mathfrak{S}$.

Theorem B generalizes some previously known results. In the case of mapping class groups: [DT15] proved equivalence of (1) and (3); equivalence of (2) and (3) follows from the distance formula; moreover, [KL08, Ham] yield that these conditions are also equivalent to convex cocompactness in the sense of [FM02]. The case of right-angled Artin groups was studied in [KMT17], where they prove equivalence of (1) and (3).

Section 6 contains a more general version of Theorem B, as well as further applications, including Theorem 6.6, which concerns the Morse boundary of hierarchically hyperbolic groups and proves that all hierarchically hyperbolic groups have finite stable asymptotic dimension.

On purely loxodromic subgroups. In the mapping class group setting [BBKL16] proved that the conditions in Theorem B are also equivalent to being undistorted and purely pseudo-Anosov. Similarly, in the right-angled Artin group setting, it was proven in [KMT17] that (1) and (3) are each equivalent to being purely loxodromic.

Subgroups of right-angled Coxeter groups all of whose elements act loxodromically on the contact graph were studied in the recent preprint [Tra, Theorem 1.4], who proved that property is equivalent to (3). Since there often exist elements in a right-angled Coxeter group which do not act loxodromically on the contact graph, his condition is not equivalent to (1); it is the ability to change the hierarchically hyperbolic structure as we do in Theorem 3.7, discussed below, which allows us to prove our more general result which characterizes *all* stable subgroups, not just the ones acting loxodromically on the contact graph.

Mapping class groups and right-angled Artin groups have the property that in their standard hierarchically hyperbolic structure they admit a universal acylindrical action on $\mathcal{C}S$, where $\mathcal{C}S$ is the hyperbolic space associated to the $\sqsubseteq$–maximal domain S. On the other hand, right-angled Coxeter groups often don't admit universal acylindrical actions on $\mathcal{C}S$ in their standard structure. Accordingly, we believe the following questions are interesting. The first item would generalize the situation in the mapping class group as established in [BBKL16], and the second item for right-angled Artin groups would generalize results proven in [KMT17], and for right-angled Coxeter groups would generalize results in [Tra]. If the second item is true for the mapping class group, this would resolve a question of Farb–Mosher [FM02]. See also [ADT17, Question 1].

Question C. *Let $(G, \mathfrak{S})$ be a hierarchically hyperbolic group which admits a universal acylindrical action on $\mathcal{C}S$, where S is the $\sqsubseteq$–maximal element in $\mathfrak{S}$. Let H be a finitely generated subgroup of G.*

- *Are the conditions in Theorem B also equivalent to: H is undistorted and acts purely loxodromically on $\mathcal{C}S$?*
- *Under what hypotheses on $(G, \mathfrak{S})$, are the conditions in Theorem B also equivalent to: H acts purely loxodromically on $\mathcal{C}S$?*

Note that in the context of Question C, an element acts loxodromically on $\mathcal{C}S$ if and only if it has positive translation length. This holds since the action is acylindrical and thus each element either acts elliptically or loxodromically.

In an early version of this paper, we asked if the second part of Question C held for all hierarchically hyperbolic groups. In the general hierarchically hyperbolic setting, however, the undistorted hypothesis is necessary, as pointed out to us by Anthony Genevois with the following example. The necessity is shown by Brady's example of a torsion-free hyperbolic group with a finitely presented subgroup which is not hyperbolic [Bra99]. This subgroup is torsion-free and thus purely loxodromic. But, a subgroup of a hyperbolic group is stable if and only if it is quasiconvex. Thus, since this subgroup is not quasiconvex, we see that being purely loxodromic is strictly weaker than the conditions of Theorem B.

New hierarchically hyperbolic structures. In order to establish the above results, we provide some new structural theorems about hierarchically hyperbolic spaces.

One of the key technical innovations in this paper is provided in Section 3. There we prove Theorem 3.7 which allows us to modify a given hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{S})$ by removing $\mathcal{C}U$ for some $U \in \mathfrak{S}$ and, in their place, enlarging the space $\mathcal{C}S$. For instance, this is how we construct the space on which a hierarchically hyperbolic group has its largest acylindrical action.

Another important tool is Theorem 4.4 which provides a simple characterization of contracting geodesics in a hierarchically hyperbolic space.

The following is a restatement of that result in the case of groups:

Theorem D (Characterization of contracting quasigeodesics)**.** *Let G be a hierarchically hyperbolic group. For any $D > 0$ and $K \geqslant 1$ there exists a $D' > 0$ depending only on D and G such that the following holds for every (K, K)–quasigeodesic $\gamma \subset \mathcal{X}$: the quasigeodesic γ has D–bounded projections if and only if γ is D'–contracting.*

Since the presence of a contracting geodesic implies the group has at least quadratic divergence, an immediate consequence of Theorem D is that any hierarchically hyperbolic group has quadratic divergence whenever $\mathcal{X}$ projects to an infinite diameter subset of $\mathcal{C}S$.

As a sample application of Theorem D and using work of Taylor–Tiozzo [TT16], we prove the following in Section 6.4 as Theorem 6.8.

Theorem E (Random subgroups are stable)**.** *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS for which $\mathcal{C}S$ has infinite diameter, where S is the $\sqsubseteq$–maximal element, and consider $G < \mathrm{Aut}(\mathcal{X}, \mathfrak{S})$ which acts properly and cocompactly on $\mathcal{X}$. Then any k–generated random subgroup of G stably embeds in $\mathcal{X}$ via the orbit map.*

We note that one immediate consequence of this result is a new proof of a theorem of Maher–Sisto: any random subgroup of a hierarchically hyperbolic group which is not the direct product of two infinite groups is stable [MS19]. The mapping class group and right-angled Artin groups cases of this result were first established in [TT16].

Finally, at the end of the paper we discuss a technical condition on hierarchically hyperbolic structures, called having *clean containers*. While in Proposition 7.2 this hypothesis is shown to hold for many groups, it does not hold in all cases. This condition was used in earlier versions of this paper in which it was assumed for the proof of Theorem 3.7, and then the general result was bootstrapped from there. In light of Theorem A.1 in the Appendix, this property is no longer required for this paper. We keep the contents of this section in the paper nonetheless, since they have found independent interest and already been used elsewhere, e.g., [BR, HS16, Rus19], as well as in several papers in progress.

Acknowledgments. The authors thank Mark Hagen and Alessandro Sisto for lively conversations about hierarchical hyperbolicity. The authors were supported in part by NSF grant DMS-1440140 while at the Mathematical Sciences Research Institute in Berkeley during Fall 2016 program in Geometric Group Theory. Abbott was supported by the NSF RTG award DMS-1502553 and NSF award DMS-1803368, Behrstock was supported by NSF award DMS-1710890, and Durham was supported by NSF RTG award DMS-1045119 and NSF award DMS-1906487. Behrstock thanks Chris Leininger for an interesting conversation which led to the formulation of Question C. We thank Daniel Berlyne, Ivan Levcovitz, Jacob Russell, Davide Spriano, and the anonymous referee for helpful feedback. We thank Anthony Genevois for resolving a question asked in an early version of this article.

2. Background

We begin by recalling some preliminary notions about metric spaces, maps between them, and group actions. Given metric spaces X, Y, we use d_X, d_Y to denote the distance functions in X, Y, respectively. A map $f \colon X \to Y$ is *K–Lipschitz* if there exists a constant $K \geqslant 1$ such that for every $x, y \in X$, $d_X(x,y) \leqslant K d_Y(g(x), g(y))$; it is *$(K,C)$–coarsely Lipschitz* if $d_X(x,y) \leqslant K d_X(x,y) + C$. The map f is a *(K,C)–quasi-isometric embedding* if there exist constants $K \geqslant 1$ and $C \geqslant 0$ such that for all $x, y \in X$,

$$\frac{1}{K} d_X(x,y) - C \leqslant d_Y(f(x), f(y)) \leqslant K d_X(x,y) + C.$$

If, in addition, Y is contained in the C–neighborhood of $f(X)$, then f is a *(K,C)–quasi-isometry*. For any interval $I \subseteq \mathbb{R}$, the image of an isometric embedding $I \to X$ is a *geodesic* and the image of a (K,C)–quasi-isometric embedding $I \to X$ is a *(K,C)–quasigeodesic*.

If any two points in X can be connected by a (K,C)–quasigeodesic, then we say X is a *(K,C)–quasigeodesic space*. If $K = C$, we may simply say that X is a *K–quasigeodesic space*. A subspace $Z \subseteq X$ is *K–quasi-convex* if there exists a constant $K \geqslant 0$ such that any geodesic in X connecting points in Z is contained in the K–neighborhood of Z. For all of the above notions, if the particular constants K, C are not important, we may drop them and simply say, for example, that a map is a quasi-isometry.

Throughout this paper, we will assume that all group actions are by isometries. The action of a group G on a metric space X, which we denote by $G \curvearrowright X$, is *proper* if for every bounded subset $B \subseteq X$, the set $\{g \in G \mid gB \cap B \neq \varnothing\}$ is finite. The action is *cobounded* (respectively, *cocompact*) if there exists a bounded (respectively, compact) subset $B \subseteq X$ such that $X = \cup_{g \in G} gB$. If a group G acts on metric spaces X and Y, we say a map $f\colon X \to Y$ is G*–equivariant* if for every $x \in X$, $f(gx) = gf(x)$. A *quasi-action* of G on X associates to each $g \in G$ a quasi-isometry $A_g\colon x \to gx$ of X with uniform quasi-isometry constants, such that $A_g \circ A_h$ is within uniformly bounded distance of A_{gh}.

2.1. **Hierarchically hyperbolic spaces.** In this section we recall the basic definitions and properties of hierarchically hyperbolic spaces as introduced in [BHS17b, BHS19].

Definition 2.1 (Hierarchically hyperbolic space)**.** A q–quasigeodesic space $(\mathcal{X}, \mathsf{d}_{\mathcal{X}})$ is said to be *hierarchically hyperbolic* if there exists $\delta \geqslant 0$, an index set $\mathfrak{S}$, and a set $\{\mathcal{C}W \mid W \in \mathfrak{S}\}$ of δ–hyperbolic spaces $(\mathcal{C}U, \mathsf{d}_U)$, such that the following conditions are satisfied:

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

(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 which is larger than all other elements; when $V \sqsubseteq W$, we say V is *nested* in W. For each $W \in \mathfrak{S}$, we denote by $\mathfrak{S}_W$ the set of $V \in \mathfrak{S}$ such that $V \sqsubseteq W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \subsetneq W$ there is a specified subset $\rho^V_W \subset \mathcal{C}W$ with $\mathsf{diam}_{\mathcal{C}W}(\rho^V_W) \leqslant \xi$. There is also a *projection* $\rho^W_V : \mathcal{C}W \to 2^{\mathcal{C}V}$.

(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 $\{V \in \mathfrak{S}_T \mid V \perp U\} \neq \emptyset$, there exists $W \in \mathfrak{S}_T - \{T\}$, so that whenever $V \perp U$ and $V \sqsubseteq T$, we have $V \sqsubseteq W$; we say W is a *container* associated with $T \in \mathfrak{S}$ and $U \in \mathfrak{S}_T$. 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 $\kappa_0 \geqslant 0$ such that if $V \pitchfork W$, then there are sets $\rho^V_W \subseteq \mathcal{C}W$ and $\rho^W_V \subseteq \mathcal{C}V$ each of diameter at most ξ and satisfying:
$$\min\left\{\mathsf{d}_W(\pi_W(x), \rho^V_W), \mathsf{d}_V(\pi_V(x), \rho^W_V)\right\} \leqslant \kappa_0$$
for all $x \in \mathcal{X}$.

For $V, W \in \mathfrak{S}$ satisfying $V \sqsubseteq W$ and for all $x \in \mathcal{X}$, we have:
$$\min\left\{\mathsf{d}_W(\pi_W(x), \rho^V_W), \mathsf{diam}_{\mathcal{C}V}(\pi_V(x) \cup \rho^W_V(\pi_W(x)))\right\} \leqslant \kappa_0.$$

Finally, if $U \sqsubseteq V$, then $\mathsf{d}_W(\rho^U_W, \rho^V_W) \leqslant \kappa_0$ whenever $W \in \mathfrak{S}$ satisfies either $V \subsetneq W$ or $V \pitchfork W$ and $W \not\perp U$.

(5) **(Finite complexity.)** There exists $n \geqslant 0$, 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 $\lambda \geqslant 1$ and $E \geqslant \max\{\xi, \kappa_0\}$ such that the following holds. Let $W \in \mathfrak{S}$ and let $x, x' \in \mathcal{X}$. Let $N = \lambda \mathsf{d}_W(\pi_W(x), \pi_W(x')) + \lambda$. Then there exists $\{T_i\}_{i=1,\dots,\lfloor N \rfloor} \subseteq \mathfrak{S}_W - \{W\}$ such that for all $T \in \mathfrak{S}_W - \{W\}$, either $T \in \mathfrak{S}_{T_i}$ for some i, or $\mathsf{d}_T(\pi_T(x), \pi_T(x')) < E$. Also, $\mathsf{d}_W(\pi_W(x), \rho^{T_i}_W) \leqslant N$ 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 $\mathcal{C}W$, either $\mathsf{diam}_{\mathcal{C}V}(\rho^W_V(\gamma)) \leqslant E$ or $\gamma \cap \mathcal{N}_E(\rho^V_W) \neq \emptyset$.

(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 $p_j \in \pi_{V_j}(\mathcal{X}) \subseteq \mathcal{C}V_j$. Then there exists $x \in \mathcal{X}$ so that:
   - $\mathsf{d}_{V_j}(x, p_j) \leqslant \alpha$ for all j,

- for each j and each $V \in \mathfrak{S}$ with $V_j \subsetneq V$, we have $\mathsf{d}_V(x, \rho_V^{V_j}) \leqslant \alpha$, and
- if $W \pitchfork V_j$ for some j, then $\mathsf{d}_W(x, \rho_W^{V_j}) \leqslant \alpha$.

(9) **(Uniqueness.)** For each $\kappa \geqslant 0$, there exists $\theta_u = \theta_u(\kappa)$ such that if $x, y \in \mathcal{X}$ and $\mathsf{d}(x,y) \geqslant \theta_u$, then there exists $V \in \mathfrak{S}$ such that $\mathsf{d}_V(x,y) \geqslant \kappa$.

Notation 2.2. Note that below we will often abuse notation by simply writing $(\mathcal{X}, \mathfrak{S})$ or $\mathfrak{S}$ to refer to the entire package of an hierarchically hyperbolic structure, including all the associated spaces, projections, and relations given by the above definition.

Notation 2.3. When writing distances in $\mathcal{C}U$ for some $U \in \mathfrak{S}$, we often simplify the notation slightly by suppressing the projection map π_U, i.e., given $x, y \in \mathcal{X}$ and $p \in \mathcal{C}U$ we write $\mathsf{d}_U(x,y)$ for $\mathsf{d}_U(\pi_U(x), \pi_U(y))$ and $\mathsf{d}_U(x,p)$ for $\mathsf{d}_U(\pi_U(x), p)$. Note that when we measure distance between a pair of sets (typically both of bounded diameter) we are taking the minimum distance between the two sets. For distance/diameter, if the space in which the measurement is being made is not clear from the context, we will denote it by a subscript. Given $A \subset \mathcal{X}$ and $U \in \mathfrak{S}$ we let $\pi_U(A)$ denote $\cup_{a \in A} \pi_U(a)$.

Remark 2.4. In the setting of hierarchically hyperbolic spaces, we often encounter maps which are well-defined only up to uniformly bounded error, in the following sense. Given a map $f \colon X \to Y$ between quasi-geodesic spaces X, Y, there may be multiple possible points in Y that one could define as $f(x)$ for a particular $x \in X$. If the diameter of such possible points $f(x)$ is uniformly bounded in Y over all $x \in X$, then we say that the map is *coarsely well-defined*, since we could arbitrarily make a choice for each $f(x)$ and the map would be well-defined up to uniformly bounded error. For example, ρ_V^U gives a coarsely well-defined map $\mathcal{C}U \to \mathcal{C}V$.

An important consequence of being a hierarchically hyperbolic space is the following distance formula, which relates distances in $\mathcal{X}$ to distances in the hyperbolic spaces $\mathcal{C}U$ for $U \in \mathfrak{S}$. The notation $\{\!\{x\}\!\}_s$ means include x in the sum if and only if $x > s$.

Theorem 2.5 (Distance formula for HHS; [BHS19])**.** *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then there exists s_0 such that for all $s \geqslant s_0$, there exist C, K so that for all $x, y \in \mathcal{X}$,*

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

We now recall an important construction of subspaces in a hierarchically hyperbolic space called *standard product regions* introduced in [BHS17b, Section 13] and studied further in [BHS19]. First we define a *consistent tuple*, which will be used to define the two factors in the product space.

Definition 2.6 (Consistent tuple)**.** Fix $\kappa \geqslant 0$, and let $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{\mathcal{C}U}$ be a tuple such that for each $U \in \mathfrak{S}$, the coordinate b_U is a subset of $\mathcal{C}U$ with $\mathsf{diam}_{\mathcal{C}U}(b_U) \leqslant \kappa$. The tuple $\vec{b}$ is *κ–admissible* if $\mathsf{d}_U(b_U, \pi_U(\mathcal{X})) \leqslant \kappa$ for all $U \in \mathfrak{S}$. The κ–admissible tuple $\vec{b}$ is *κ–consistent* if, whenever $V \pitchfork W$,

$$\min\left\{\mathsf{d}_W(b_W, \rho_W^V), \mathsf{d}_V(b_V, \rho_V^W)\right\} \leqslant \kappa$$

and whenever $V \sqsubseteq W$,

$$\min\left\{\mathsf{d}_W(b_W, \rho_W^V), \mathsf{diam}_{\mathcal{C}V}(b_V \cup \rho_V^W(b_W))\right\} \leqslant \kappa.$$

Definition 2.7 (Nested partial tuple $(\mathbf{F}_U)$)**.** Recall $\mathfrak{S}_U = \{V \in \mathfrak{S} \mid V \sqsubseteq U\}$. Fix $\kappa \geqslant \kappa_0$ and let $\mathbf{F}_U$ be the set of κ–consistent tuples in $\prod_{V \in \mathfrak{S}_U} 2^{\mathcal{C}V}$.

Definition 2.8 (Orthogonal partial tuple ($\mathbf{E}_U$))**.** Let $\mathfrak{S}_U^{\perp} = \{V \in \mathfrak{S} \mid V \perp U\} \cup \{A\}$, where A is a $\sqsubseteq$–minimal element W such that $V \sqsubseteq W$ for all $V \perp U$ (note that A exists by the container axiom for an HHS, i.e., Definition 2.1.(3)). Fix $\kappa \geqslant \kappa_0$, let $\mathbf{E}_U$ be the set of κ–consistent tuples in $\prod_{V \in \mathfrak{S}_U^{\perp} - \{A\}} 2^{\mathcal{C}V}$.

Definition 2.9 (Product regions in $\mathcal{X}$)**.** Given $\mathcal{X}$ and $U \in \mathfrak{S}$, there is a coarsely well-defined map $\phi_U \colon \mathbf{F}_U \times \mathbf{E}_U \to \mathcal{X}$ which restricts to coarsely well-defined maps $\phi^{\sqsubseteq}, \phi^{\perp} \colon \mathbf{F}_U, \mathbf{E}_U \to \mathcal{X}$. Indeed, for each $(\vec{a}, \vec{b}) \in \mathbf{F}_U \times \mathbf{E}_U$, and each $V \in \mathfrak{S}$, the projection $\pi_V(\phi_U(\vec{a}, \vec{b}))$ is defined as follows. If $V \sqsubseteq U$, then $\pi_V(\phi_U(\vec{a}, \vec{b})) = a_V$. If $V \perp U$, then $\pi_V(\phi_U(\vec{a}, \vec{b})) = b_V$. If $V \pitchfork U$, then $\pi_V(\phi_U(\vec{a}, \vec{b})) = \rho^U_V$. Finally, if $U \sqsubseteq V$, and $U \neq V$, let $\pi_V(\phi_U(\vec{a}, \vec{b})) = \rho^U_V$. The tuple $(\pi_V(\phi_U(\vec{a}, \vec{b})))_{V \in \mathfrak{S}} \in \prod_{V \in \mathfrak{S}} 2^{\mathcal{C}V}$ is κ–consistent (see [BHS19, Construction 5.10]), and therefore [BHS19, Theorem 3.1] provides a point $x \in \mathcal{X}$ such that $\mathsf{d}_W(\pi_W(x), \pi_W(\phi_U(\vec{a}, \vec{b}))) \leqslant \theta_e$ for all $W \in \mathfrak{S}$. Moreover, the point x is *coarsely unique* in the sense that the set of all x which satisfy $\mathsf{d}_W(\pi_W(x), \pi_W(\phi_U(\vec{a}, \vec{b}))) \leqslant \theta_e$ for each $W \in \mathfrak{S}$ has diameter at most θ_u in $\mathcal{X}$. We define $\phi_U(\vec{a}, \vec{b}) = x$; the coarse uniqueness of x shows that this map is coarsely well-defined. Fixing any $e \in \mathbf{E}_U$ yields a map $\phi_U^{\sqsubseteq} \colon \mathbf{F}_U \times \{e\} \to \mathcal{X}$, and $\phi^{\perp}$ is defined analogously. We refer to $\mathbf{F}_U \times \mathbf{E}_U$ as a *product region*, which we denote $\mathbf{P}_U$.

We often abuse notation slightly and use the notation $\mathbf{E}_U, \mathbf{F}_U$, and $\mathbf{P}_U$ to refer to the image in $\mathcal{X}$ of the associated set. In [BHS19, Construction 5.10] it is proven that these standard product regions have the property that they are "hierarchically quasiconvex subsets" of $\mathcal{X}$. We leave out the definition of hierarchically quasiconvexity, because its only use here is that product regions have "gate maps," as given by the following in [BHS19, Lemma 5.5]:

Lemma 2.10 (Existence of coarse gates; [BHS19, Lemma 5.5])**.** *If $\mathcal{Y} \subseteq \mathcal{X}$ is k–hierarchically quasiconvex and non-empty, then there exists a gate map for $\mathcal{Y}$, i.e., for each $x \in \mathcal{X}$ there exists $\mathfrak{g}(x) \in \mathcal{Y}$ such that for all $V \in \mathfrak{S}$, the set $\pi_V(\mathfrak{g}(x))$ (uniformly) coarsely coincides with the projection of $\pi_V(x)$ to the $k(0)$–quasiconvex set $\pi_V(\mathcal{Y})$. The point $\mathfrak{g}(x) \in \mathcal{Y}$ is called the* gate of x in $\mathcal{Y}$.

Remark 2.11 (Surjectivity of projections)**.** As one can always change the hierarchical structure so that the projection maps are coarsely surjective [BHS19, Remark 1.3], we will assume that $\mathfrak{S}$ is such a structure. That is, for each $U \in \mathfrak{S}$, if π_U is not surjective, then we identify $\mathcal{C}U$ with $\pi_U(\mathcal{X})$.

We also need the notion of a hierarchy path, whose existence was proven in [BHS19, Theorem 4.4] (although we use the word *path*, since they are quasi-geodesics, typically we consider them as discrete sequences of points):

Definition 2.12. For $R \geqslant 1$, a path γ in $\mathcal{X}$ is a *R–hierarchy path* if

(1) γ is a (R, R)–quasigeodesic,
(2) for each $W \in \mathfrak{S}$, $\pi_W \circ \gamma$ is an unparametrized (R, R)–quasigeodesic. An unbounded hierarchy path $[0, \infty) \to \mathcal{X}$ is a *hierarchy ray.*

We call a domain *relevant* to a pair of points, if the projections to that domain are larger than some fixed (although possibly unspecified) constant depending only on the hierarchically hyperbolic structure. We say a domain is *relevant* for a particular quasi-geodesic if it is relevant for the endpoints of that quasi-geodesic.

Proposition 2.13 ([BHS19, Proposition 5.17])**.** *There exists $\nu \geqslant 0$ such that for all $x, y \in \mathcal{X}$, all $V \in \mathfrak{S}$ with V relevant for (x, y), and all D–hierarchy paths γ joining x to y, there is a subpath α of γ with the following properties:*

(1) $\alpha \subset \mathcal{N}_\nu(\mathbf{P}_V)$;

(2) $\pi_U|_\gamma$ *is coarsely constant on* $\gamma - \alpha$ *for all* $U \in \mathfrak{S}_V \cup \mathfrak{S}_V^\perp$*, i.e., it is a uniformly bounded distance from a constant map.*

Remark 2.14. Let $x, y \in \mathcal{X}$, and suppose V is relevant for (x, y). As $\mathbf{F}_V$ and $\mathbf{E}_V$ consist of κ–consistent tuples (for a fixed κ) and $\phi_V : \mathbf{F}_V \times \mathbf{E}_V \to \mathcal{X}$ is only coarsely well-defined, by appropriately increasing κ to accomodate for the chosen constant ν in Proposition 2.13, we may assume that α is actually a subset of $\mathbf{P}_V$.

It is often convenient to work with equivariant hierarchically hyperbolic structures, we now recall the relevant structures for doing so. For details see [BHS19].

Definition 2.15 (Hierarchically hyperbolic groups)**.** Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. An *automorphism* of $(\mathcal{X}, \mathfrak{S})$ consists of a map $g : \mathcal{X} \to \mathcal{X}$, together with a bijection $g^\diamondsuit : \mathfrak{S} \to \mathfrak{S}$ and, for each $U \in \mathfrak{S}$, an isometry $g^*(U) : \mathcal{C}U \to \mathcal{C}(g^\diamondsuit(U))$ so that the following diagrams commute up to uniformly bounded error whenever the maps in question are defined (i.e., when U, V are not orthogonal):

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

and

$$\begin{array}{ccc} \mathcal{C}U & \xrightarrow{\ g^*(U)\ } & \mathcal{C}(g^\diamondsuit(U)) \\ {\scriptstyle \rho^U_V}\downarrow & & \downarrow{\scriptstyle \rho^{g^\diamondsuit(U)}_{g^\diamondsuit(V)}} \\ \mathcal{C}V & \xrightarrow{\ g^*(V)\ } & \mathcal{C}(g^\diamondsuit(V)) \end{array}$$

Two automorphisms f, f' are *equivalent* if $f^\diamond = (f')^\diamond$ and for all $U \in \mathfrak{S}$ we have $\phi_U = \phi'_U$. The set of all such equivalence classes forms the *automorphism group* of $(\mathcal{X}, \mathfrak{S})$, denoted $\mathrm{Aut}(\mathcal{X}, \mathfrak{S})$. A finitely generated group G is said to be a *hierarchically hyperbolic group (HHG)* if there is a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ and a group homomorphism $G \to \mathrm{Aut}(\mathcal{X}, \mathfrak{S})$ so that the induced uniform quasi-action of G on $\mathcal{X}$ is metrically proper, cobounded, and $\mathfrak{S}$ contains finitely many G–orbits. Note that when G is a hyperbolic group then, with respect to any word metric, it inherits a hierarchically hyperbolic structure.

2.2. **Acylindrical actions.** We recall the basic definitions related to acylindrical actions; the canonical references are [Bow08] and [Osi16]. We also discuss a partial order on these actions which was recently introduced in [ABO19].

Definition 2.16 (Acylindrical)**.** The action of a group G on a metric space X is *acylindrical* if for any $\varepsilon > 0$ there exist $R, N \geqslant 0$ such that for all $x, y \in X$ with $\mathsf{d}(x, y) \geqslant R$,

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

Recall that given a group G acting on a hyperbolic metric space X, an element $g \in G$ is *loxodromic* if the map $\mathbb{Z} \to X$ defined by $n \mapsto g^n x$ is a quasi-isometric embedding for some (equivalently any) $x \in X$. However, an element of G may be loxodromic for some actions and not for others. Consider, for example, the free group on two generators acting on its Cayley graph and acting on the Bass-Serre tree associated to the splitting $\mathbb{F}_2 \simeq \langle x \rangle * \langle y \rangle$. In the former action, every non-trivial element is loxodromic, while in the latter action, no powers of x and y are loxodromic.

Definition 2.17 (Generalized loxodromic)**.** An element of a group G is called *generalized loxodromic* if it is loxodromic for some acylindrical action of G on a hyperbolic space.

Remark 2.18. By [Osi16, Theorem 1.1], every acylindrical action of a group on a hyperbolic space either has bounded orbits or contains a loxodromic element. By [Osi16, Theorem 1.4.(L4)] and Sisto [Sis16, Theorem 1], every generalized loxodromic element is *Morse*, i.e., every quasi-geodesic with endpoints on the axis of the element lies uniformly close to that axis (see Definition 2.22). Therefore, if a group H does not contain any Morse elements, it does not contain any generalized loxodromics, and thus H must have bounded orbits in every acylindrical action on a hyperbolic space. This is the case when, for example, H is a non-trivial direct product, that is, a direct product of two infinite groups.

Definition 2.19 (Universal acylindrical action)**.** An acylindrical action of a group on a hyperbolic space is a *universal acylindrical action* if every generalized loxodromic element is loxodromic. Such an action is sometimes called a *loxodromically universal action*.

Notice that if every acylindrical action of a group G on a hyperbolic space has bounded orbits, then G does not contain any generalized loxodromic elements, and the action of G on a point (which is acylindrical) is a universal acylindrical action.

The following notions are discussed in detail in [ABO19]. We give a brief overview here. Fix a group G. Given a (possibly infinite) generating set X of G, let $|\cdot|_X$ denote the word metric with respect to X, and let $\Gamma(G,X)$ be the Cayley graph of Γ with respect to the generating set X. Given two generating sets X and Y, we say X is *dominated* by Y and write $X \preceq Y$ if

$$\sup_{y\in Y} |y|_X < \infty.$$

Note that when $X \preceq Y$, the action $G \curvearrowright \Gamma(G,Y)$ provides more information about the group than $G \curvearrowright \Gamma(G,X)$, and so, in some sense, is a "larger" action. The two generating sets X and Y are equivalent if $X \preceq Y$ and $Y \preceq X$; when this happens we write $X \sim Y$.

Let $\mathcal{AH}(G)$ be the set of equivalence classes of generating sets X of G such that $\Gamma(G,X)$ is hyperbolic and the action $G \curvearrowright \Gamma(G,X)$ is acylindrical. We denote the equivalence class of X by $[X]$. The preorder $\preceq$ induces an order on $\mathcal{AH}(G)$, which we also denote $\preceq$.

Definition 2.20 (Largest)**.** We say an equivalence class of generating sets is *largest* if it is the largest element in $\mathcal{AH}(G)$ under this ordering.

Given a cobounded acylindrical action of G on a hyperbolic space S, a Milnor–Schwartz argument gives a (possibly infinite) generating set Y of G such that there is a G–equivariant quasi-isometry between $G \curvearrowright S$ and $G \curvearrowright \Gamma(G,Y)$. By a slight abuse of language, we will say that a particular cobounded acylindrical action $G \curvearrowright S$ on a hyperbolic space is largest, when, more precisely, it is the equivalence class of the generating set associated to this action through the above correspondence, $[Y]$, that is the largest element in $\mathcal{AH}(G)$.

Remark 2.21. By definition, every largest acylindrical action is a universal acylindrical action. To see this, notice that if $[X] \preceq [Y]$, then the set of loxodromic elements in $G \curvearrowright \Gamma(G,X)$ must be a subset of the set of loxodromic elements in $G \curvearrowright \Gamma(G,Y)$.

2.3. **Stability.** Stability is strong coarse convexity property which generalizes quasiconvexity in hyperbolic spaces and convex cocompactness in mapping class groups [DT15]. In the general context of metric spaces, it is essentially the familiar Morse property generalized to subspaces, so we begin there.

Definition 2.22 (Morse/stable quasigeodesic)**.** Let X be a metric space. A quasigeodesic $\gamma \subset X$ is called *Morse* (or *stable*) if there exists a function $N\colon \mathbb{R}^2_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ such that if q is a (K,C)–quasigeodesic in X with endpoints on γ, then

$$q \subset \mathcal{N}_{N(K,C)}(\gamma).$$

We call N the *stability gauge* for γ and say γ is N–stable if we want to record the constants.

We can now define a notion of stable embedding of one metric space in another which is equivalent to the one introduced by Durham and Taylor [DT15]:

Definition 2.23 (Stable embedding)**.** We say a quasi-isometric embedding $f\colon X \to Y$ between quasigeodesic metric spaces is a *stable embedding* if there exists a stability gauge N such that for any quasigeodesic constants K, C and any (K, C)–quasigeodesic $\gamma \subset X$, then $f(\gamma)$ is an $N(K, C)$–stable quasigeodesic in Y. We say a subset $X \subseteq Y$ is *stable* if it is undistorted and the inclusion map $i\colon X \to Y$ is a stable embedding.

The following generalizes the notion of a Morse quasigeodesic to subgroups:

Definition 2.24 (Subgroup stability)**.** Let H be subgroup of a finitely generated group G. We say H is a *stable subgroup* of G if some (equivalently, any) orbit map of H into some (any) Cayley graph (with respect to a finite generating set) of G is a stable embedding.

If for some $h \in G$, $H = \langle h\rangle$ is stable, then we call h *stable.* Such elements are often called *Morse elements.*

Stability of a subset is preserved under quasi-isometries. Note that stable subgroups are undistorted in their ambient groups and, moreover, they are quasiconvex with respect to any choice of finite generating set for the ambient group.

3. Altering the hierarchically hyperbolic structure

The goal of this section is to prove that any hierarchically hyperbolic space satisfying a technical assumption—the *bounded domain dichotomy*—admits a hierarchically hyperbolic structure with unbounded products, i.e., every non-trivial product region in the ambient space has unbounded factors; see Theorem 3.7 below.

In particular, this establishes that all hierarchically hyperbolic groups admit a hierarchically hyperbolic group structure with unbounded products. It is for this reason that our complete characterization of the contracting property in spaces with unbounded products in Section 4 yields a characterization of the contracting property for all hierarchically hyperbolic groups, as stated in Theorem D.

3.1. **Unbounded products.** Fix a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$.

Let $M > 0$ and let $\mathfrak{S}^M \subset \mathfrak{S}$ be the set of domains $U \in \mathfrak{S}$ such that there exists $V \in \mathfrak{S}$ and $W \in \mathfrak{S}_V^\perp$ satisfying: $U \sqsubseteq V$, $\mathsf{diam}(\mathcal{C}V) > M$, and $\mathsf{diam}(\mathcal{C}W) > M$.

Recall that a set of domains $\mathfrak{U} \subset \mathfrak{S}$ is *closed under nesting* if whenever $U \in \mathfrak{U}$ and $V \sqsubseteq U$, then $V \in \mathfrak{U}$.

Lemma 3.1. *For any $M > 0$, the set $\mathfrak{S}^M$ is closed under nesting.*

Proof. Let $U \in \mathfrak{S}^M$ and $V \sqsubseteq U$. By definition of $U \in \mathfrak{S}^M$, there exists $Z \in \mathfrak{S}^M$ with $U \sqsubseteq Z$ and satisfying: $\mathsf{diam}(\mathcal{C}Z) > M$ and there exists $W \in \mathfrak{S}_Z^\perp$ such that $\mathsf{diam}(\mathcal{C}W) > M$. Since $V \sqsubseteq Z$, it follows that $V \in \mathfrak{S}^M$, as desired. □

Definition 3.2 (Bounded domain dichotomy)**.** We say $(\mathcal{X}, \mathfrak{S})$ has the *M–bounded domain dichotomy* if there exists $M > 0$ such that any $U \in \mathfrak{S}$ with $\mathsf{diam}(\mathcal{C}U) > M$ satisfies $\mathsf{diam}(\mathcal{C}U) = \infty$. If the value of M is not important, we simply refer to the *bounded domain dichotomy.*

Recall that for every hierarchically hyperbolic group $(G, \mathfrak{S})$, the set of domains $\mathfrak{S}$ contains finitely many G–orbits and each $g \in G$ induces an isometry $\mathcal{C}U \to \mathcal{C}(g^\diamond(U))$ for each $U \in \mathfrak{S}$ (see Definition 2.15). It thus follows that every hierarchically hyperbolic group has the bounded domain dichotomy. (Also, note that this property implies the space is "asymphoric" as defined in [BHS17c].)

Definition 3.3 (Unbounded products)**.** We say that a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ has *unbounded products* if it has the bounded domain dichotomy and the property that if $U \in \mathfrak{S} - \{S\}$ has $\mathsf{diam}(\mathbf{F}_U) = \infty$, then $\mathsf{diam}(\mathbf{E}_U) = \infty$.

3.2. **Almost hierarchically hyperbolic spaces.** In this section we introduce a tool for verifying a space is hierarchically hyperbolic.

The following is a weaker version of the orthogonality axiom:

(3′) **(Bounded pairwise 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 $V \perp W$. Moreover, if $V \perp W$, then V, W are not $\sqsubseteq$–comparable. Finally, the cardinality of any collection of pairwise orthogonal domains is uniformly bounded by ξ.

By [BHS19, Lemma 2.1], the orthogonality axiom (Definition 2.1, (3)) for an hierarchically hyperbolic structure implies axiom (3′). However, the converse does not hold; that is, the last condition of (3′) does not directly imply the container statement in (3), and thus this is an *a priori* strictly weaker assumption. However, as is proven in the appendix in Theorem A.1, this weakened version of the axiom is enough to produce a hierarchically hyperbolic structre.

We now introduce the notion of an *almost hierarchically hyperbolic space*:

Definition 3.4 (Almost HHS)**.** If $(\mathcal{X}, \mathfrak{S})$ satisfies all axioms of a hierarchically hyperbolic space except (3) and additionally satisfies axiom (3′), then $(\mathcal{X}, \mathfrak{S})$ is an *almost hierarchically hyperbolic space.*

In the appendix, Berlyne and Russell prove Theorem A.1, establishing that if a space is almost hierarchically hyperbolic, then the associated structure can be modified to obtain a hierarchically hyperbolic structure on the original space. This result is used in our proof of Theorem 3.7.

3.3. **A new hierarchically hyperbolic structure.** In this section we describe a new hierarchically hyperbolic structure on hierarchically hyperbolic spaces with the bounded domain dichotomy. We first describe the hyperbolic spaces that will be part of the new structure.

Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space with the M–bounded domain dichotomy. Recall that we define $\mathfrak{S}^M \subset \mathfrak{S}$ to be the set of $U \in \mathfrak{S}$ such that there exists $U \sqsubseteq V$ with $\mathsf{diam}(\mathcal{C}V) > M$ for which there exists a $W \in \mathfrak{S}_V^{\perp}$ satisfying $\mathsf{diam}(\mathcal{C}W) > M$. For each $U \in \mathfrak{S}$, define $\mathfrak{S}_U^M \subset \mathfrak{S}_U$ similarly.

Remark 3.5 (Factored spaces)**.** As defined in [BHS17a], given $(\mathcal{X}, \mathfrak{S})$ and $\mathfrak{T} \subset \mathfrak{S}$ the *factored space* $\widehat{\mathbf{F}}_{\mathfrak{T}}$ is the space obtained from $\mathcal{X}$ by coning-off each $\mathbf{F}_V \times \{e\}$ for all $V \in \mathfrak{T}$ and all $e \in \mathbf{E}_V$. Sometimes we abuse language slightly and refer to this as the factored space obtained from $\mathcal{X}$ by collapsing $\mathfrak{T}$. In particular, when S is the $\sqsubseteq$–maximal element of $\mathfrak{S}$, then $\mathcal{C}S$ can be taken to be the space $\widehat{\mathbf{F}}_{\mathfrak{S}-\{S\}}$, which is obtained from $\mathcal{X}$ by coning-off $\mathbf{F}_U \times \{e\}$ for all $U \in \mathfrak{S} - \{S\}$ and all $e \in \mathbf{E}_U$.

We often consider the case of a fixed $(\mathcal{X}, \mathfrak{S})$ and $U \in \mathfrak{S}$ and then applying this construction to the hierarchy hyperbolic structure $(\mathbf{F}_U, \mathfrak{S}_U)$. For this application, note that $\pi_U(\mathcal{X})$ is quasi-isometric to $\widehat{\mathbf{F}}_{\mathfrak{S}_U - \{U\}}$, by [BHS17a, Corollary 2.9], and thus so is $\mathcal{C}U$, by Remark 2.11.

Lemma 3.6. *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space and consider $\mathfrak{T} \subset \mathfrak{S}$ which is closed under nesting. Let γ be a hierarchy path in $(\mathcal{X}, \mathfrak{S})$. Then, the path obtained by including $\gamma \subset \mathcal{X} \subset \widehat{\mathbf{F}}_{\mathfrak{T}}$ is an unparametrized quasi-geodesic. Moreover, if for each $W \in \mathfrak{T}$ which is a relevant domain for γ and for each $e \in \mathbf{E}_W$, we modify the path through $\mathbf{F}_W \times \{e\}$ by removing all but the first and last vertex of the hierarchy path which passes through $\mathbf{F}_W \times \{e\}$, then the new path obtained, $\hat{\gamma}$ is a hierarchy path for $(\widehat{\mathbf{F}}_{\mathfrak{T}}, \mathfrak{S} - \mathfrak{T})$.*

Proof. The proof is by induction on complexity. Consider all the nest-minimal elements $\mathfrak{U} \subset \mathfrak{T}$ which are relevant for γ; by Proposition 2.13 and Remark 2.14 for each such U there is a subpath of γ which passes through a collection of slices $\mathbf{F}_U \times \{e\}$ within the product region associated to U. By [BHS19, Lemma 2.14] there is a bounded (in terms of $\mathfrak{S}$) coloring of $\mathfrak{U}$ with the property that all the domains of a given color are pairwise transverse. Starting from $(\mathcal{X}, \mathfrak{S})$, we take one color at a time, together with all the domains nested inside domains of that color, and create the factored space by coning off those domains. At each step, we obtain a new hierarchically hyperbolic space with the property that in this space the relevant domains for γ are exactly the original ones except for those in the colors we have coned off thus far. Since this path still travels monotonically through each of the relevant domains, it is an unparametrized quasi-geodesic in the new factored space. Thus the path $\hat{\gamma}$ is a parametrized quasi-geodesic and thus a hierarchy path in the new factored space (with constants depending only on the constants for the original hierarchy path). Once the colors of $\mathfrak{U}$ are exhausted, repeat one step up the nesting lattice. Since the complexity of a hierarchically hyperbolic and the coloring are both bounded, this will terminate after finitely many steps. Finally we cone off any domains in $\mathfrak{T}$ which are not relevant for γ to obtain the space $(\widehat{\mathbf{F}}_{\mathfrak{T}}, \mathfrak{S} - \mathfrak{T})$. Through this final step $\hat{\gamma}$ remains a uniform quality hierarchy path since it is still a quasigeodesic. □

The next result uses the above spaces to obtain a hierarchically hyperbolic structure with particularly nice properties from a given hierarchically hyperbolic structure.

Theorem 3.7. *Every hierarchically hyperbolic space with the bounded domain dichotomy admits a hierarchically hyperbolic structure with unbounded products.*

Proof. Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Let $\mathfrak{T}$ denote the $\sqsubseteq$–maximal element S together with the subset of $\mathfrak{S}$ consisting of all $U \in \mathfrak{S}$ with both $\mathbf{F}_U$ and $\mathbf{E}_U$ unbounded.

We begin to define our new hierarchically hyperbolic structure on $\mathcal{X}$ by taking $\mathfrak{T}$ as our index set. For each $U \in \mathfrak{T} - \{S\}$ we set the associated hyperbolic space $\mathcal{T}_U$ to be $\mathcal{C}U$. For the top-level domain, S, we obtain a hyperbolic space, $\mathcal{T}_S$, as follows. By Lemma 3.1, $\mathfrak{S}^M$ is closed under nesting and hence $\widehat{\mathcal{X}}_{\mathfrak{S}^M}$ is a hierarchically hyperbolic space. Moreover, since this hierarchically hyperbolic space has the property that no pair of orthogonal domains both have diameter larger than M, by [BHS17c, Corollary 2.16] it is hyperbolic for some constant depending only on $(\mathcal{X}, \mathfrak{S})$ and M; we call this space $\mathcal{T}_S$.

To avoid confusion, we use the notation d_S for distance in $\mathcal{T}_S$ and the notation $\mathsf{d}_{\mathcal{C}S}$ for distance in $\mathcal{C}S$.

When $U \neq S$, the projections are as defined in the original hierarchically hyperbolic space. We take the projection π_S to be the factor map $\mathcal{X} \to \mathcal{T}_S$. If $U \in \mathfrak{T}$ and $U \neq S$, then the relative projections are defined as in $(\mathcal{X}, \mathfrak{S})$. For the remaining cases the relative projections are as follows: ρ^S_V is defined to be π_V and ρ^V_S is defined to be the image of $\mathbf{F}_V$ under the factor map $\mathcal{X} \to \mathcal{T}_S$.

We now check the axioms to verify that $(\mathcal{X}, \mathfrak{T})$ is an almost hierarchically hyperbolic space (i.e., all the conditions of a hierarchically hyperbolic space except for a weakened version of the orthogonality axiom). Once these axioms have been verified, we can then invoke Theorem A.1 to conclude that the almost hierarchically hyperbolic structure $\mathfrak{T}$ can be modified to yield an actual hierarchically hyperbolic space. By construction, $(\mathcal{X}, \mathfrak{T})$ satisfies the hypothesis of Corollary A.8, and therefore the associated modified hierarchically hyperbolic structure will have unbounded products, as desired.

Projections: The only case to check is for the top-level domain S. Since π_S is a factor map, it is coarsely Lipschitz and coarsely surjective.

Nesting: The partial order and projections are given by construction. The diameter bound in the case of nesting projections is immediate from the bound from $(\mathcal{X}, \mathfrak{S})$, except in the case of ρ_S^V for $V \in \mathfrak{T}$. The bound on the diameter of ρ_S^V follows from the construction of $\mathcal{T}_S$ as a factor space and the fact that $\mathfrak{T} \subset \mathfrak{S}^M$.

Orthogonality: We now verify axiom (3′) is satisfied by this new structure. Tthe first three conditions are clear, since $\mathfrak{T} \subseteq \mathfrak{S}$ and thus they are inherited from the hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{S})$. For the last condition, any collection of pairwise orthogonal domains in $\mathfrak{T}$ is also a collections of pairwise orthogonal domains in $\mathfrak{S}$, and thus by [BHS19, Lemma 2.2] has uniformly bounded size, verifying the axiom.

Transversality and consistency: This axiom only involves domains which are not nest-maximal, and hence holds using the original constants from the hierarchically hyperbolic structure on $(\mathcal{X}, \mathfrak{S})$.

Partial realization: This axiom only involves domains which are not nest-maximal, and hence holds using the original constants from the hierarchically hyperbolic structure on $(\mathcal{X}, \mathfrak{S})$.

Finite complexity: This clearly holds by construction.

Large link axiom: Let λ and E be the constants from the large link axiom for $(\mathcal{X}, \mathfrak{S})$, let $W \in \mathfrak{T}$, and let $x, x' \in \mathcal{X}$. Consider the set $\{T_i\} \subset \mathfrak{S}_W - \{W\}$ provided by the large link axiom for $(\mathcal{X}, \mathfrak{S})$. Since $T_i \sqsubseteq W$, it follows that $\mathbf{E}_{T_i}$ is unbounded for each i. Let $T \in \mathfrak{T}_W - \{W\}$. If $\mathsf{d}_T(x, x') > E \cdot M$, it follows that $\mathbf{F}_T$ is unbounded. Furthermore, $\mathsf{d}_{\mathcal{C}T}(x, x') > E$, whence $T \sqsubseteq T_i$ for some i by the large link axiom for $(\mathcal{X}, \mathfrak{S})$. Therefore $\mathbf{F}_{T_i}$ is unbounded, and so $T_i \in \mathfrak{T}$. The result follows.

Bounded geodesic image: For all domains in $\mathfrak{T} - \{S\}$, the corresponding hyperbolic spaces are unchanged from those in the original structure and thus the axiom holds in these cases.

Hence the only case which it remains to check is when $W = S$. Suppose γ is a geodesic in $\mathcal{T}_S$, and $V \in \mathfrak{T} - \{S\}$ such that $\mathsf{diam}_{\mathcal{C}V}(\rho_V^S(\gamma)) > E$. The partial realization axiom implies that there exists a hierarchy path $\bar{\gamma} \subset \mathcal{X}$ whose end-points project under π_S to the end-points of γ. This projected path is a quasigeodesic by Lemma 3.6. Since $\mathcal{T}_S$ is hyperbolic, the projected path lies uniformly close to γ. By [BHS19, Proposition 5.17]) we can replace $\bar{\gamma}$ by an appropriate subpath for which the only relevant domains are all nested in V; thus $\bar{\gamma} \subset \mathbf{P}_V$. By definition, there is a bounded distance between ρ_S^V and $\pi_S(\mathbf{P}_V)$; thus $\pi_S(\bar{\gamma})$ (and hence γ) is a bounded distance from ρ_S^V, as needed.

Uniqueness: Let $\kappa > 0$. We can take $\theta_u' > \max\{\theta_u(\kappa), M\}$, where $\theta_u(\kappa)$ is the original constant from the uniqueness axiom for $(\mathcal{X}, \mathfrak{S})$. Then if $x, y \in \mathcal{X}$ with $\mathsf{d}(x, y) > \theta_u'$, then uniqueness for $(\mathcal{X}, \mathfrak{S})$ implies there exists $U \in \mathfrak{S}$ with $\mathsf{d}_{\mathcal{C}U}(x, y) > M$. Either $U \in \mathfrak{T}$ or $\mathsf{diam}(\mathcal{C}U) = \infty$ and $\mathbf{E}_U$ is bounded. We are done in the first case. In the second case, by construction the factor space $\hat{U}$ of $\mathbf{F}_U$ obtained by factoring $\mathfrak{T}_U$ is quasi-isometrically embedded in $\mathcal{T}_S$ and there is a 1–Lipschitz map from $\hat{U}$ to $\mathcal{C}U$. Thus the lower bound on distance in $\mathcal{C}U$ provides a lower bound on the distance in $\hat{U}$, which, in turn, provides a lower bound in $\mathcal{T}_S$, as desired. □

Corollary 3.8. *Every hierarchically hyperbolic group admits a hierarchically hyperbolic group structure with unbounded products.*

Proof. Recall that every hierarchically hyperbolic group has the bounded domain dichotomy. Accordingly, if we start with a hierarchically hyperbolic group, $(G, \mathfrak{S})$, then Theorem 3.7 yields a hierarchically hyperbolic structure with unbounded products, $(G, \mathfrak{T})$, where $\mathfrak{T}$ is the structure from the proof of Theorem 3.7 with the additional "dummy domains" added as provided at the end of that proof via Theorem A.1. It remains only to show that this is a hierarchically hyperbolic group structure. The action of G on itself, by left multiplication, is clearly metrically proper and cobounded, and thus it only remains to show that $\mathfrak{T}$ contains finitely many G–orbits. If $U \in \mathfrak{S}$ but $U \notin \mathfrak{T}$, then either $\mathbf{F}_U$ or $\mathbf{E}_U$ must be bounded.

Then for each $g \in G$, the same will be true for $\mathbf{F}_{gU}$ or $\mathbf{E}_{gU}$, which shows that $gU \notin \mathfrak{T}$. Thus $G \cdot U \subset \mathfrak{S} - \mathfrak{T}$. The now result follows from the fact that $\mathfrak{S}$ has only finitely many G–orbits and that any dummy domains added fall into only finitely many orbits, as noted in Remark A.7. □

4. Characterization of contracting geodesics

For this section, fix a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ with the bounded domain dichotomy; denote the $\sqsubseteq$–maximal element $S \in \mathfrak{S}$.

Definition 4.1 (Bounded projections)**.** Let $\mathcal{Y} \subset \mathcal{X}$ and $D > 0$. We say that $\mathcal{Y}$ has *D–bounded projections* if for every $U \in \mathfrak{S} - \{S\}$, we have $d_U(\mathcal{Y}) < D$.

Definition 4.2 (Contracting)**.** A subset γ in a metric space X is said to be *D–contracting* if there exist a map $\pi_\gamma \colon X \to \gamma \subset X$ and constants $A, D > 0$ satisfying:

1. For any $x \in \gamma$, we have $\mathsf{diam}_X(x, \pi_\gamma(x)) < D$;
2. If $x, y \in X$ with $d_X(x, y) < 1$, then $\mathsf{diam}_X(\pi_\gamma(x), \pi_\gamma(y)) < D$;
3. For all $x \in X$, if we set $R = A \cdot d(x, \gamma)$, then $\mathsf{diam}_X(\pi_\gamma(B_R(x))) \leqslant D$.

In this section, we will focus our attention to the case of Definition 4.2 where γ is a quasigeodesic. In Section 6 we will consider results about arbitary subsets with the contracting property.

We note that sometimes authors refer to any quasigeodesic satisfying (3) as *contracting*. Nonetheless, for applications one also needs to assume the coarse idempotence and coarse Lipschitz properties given by (1) and (2), so for convenience we combine them all in one property.

A useful well-known fact is stability of contracting quasigeodesics. Two different proofs of the following occur as special cases of the results [MM99, Lemma 6.1] and [Beh06, Theorem 6.5]; this explicit statement can also be found in [DT15, Section 4].

Lemma 4.3. *If γ is a D–contracting (K, K)–quasigeodesic in a metric space X, then γ is D'–stable for some D' depending only on D and K.*

The following result and argument both generalize and simplify the analogous result for mapping class groups in [Beh06].

Theorem 4.4. *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. For any $D > 0$ and $K \geqslant 1$ there exists a $D' > 0$ depending only on D and $(\mathcal{X}, \mathfrak{S})$ such that the following holds for every (K, K)–quasigeodesic $\gamma \subset \mathcal{X}$. If γ has D–bounded projections, then γ is D'–contracting. Moreover, if $(\mathcal{X}, \mathfrak{S})$ has the bounded domain dichotomy, then $\mathcal{X}$ admits a hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{T})$ with unbounded products where, additionally, we have that if γ is D–contracting, then γ has D'–bounded projections.*

Proof. First suppose that γ has D–bounded projections. It follows immediately from the definition that γ is a hierarchically quasiconvex subset of $\mathcal{X}$. Hierarchical quasiconvexity is the hypothesis necessary to apply [BHS17a, Lemma 5.5] (see Lemma 2.10), which then yields existence of a coarsely Lipschitz *gate map* $\mathfrak{g} \colon \mathcal{X} \to \gamma$, i.e., for each $x \in \mathcal{X}$, the image $\mathfrak{g}(x) \in \gamma$ has the property that for all $U \in \mathfrak{S}$ the set $\pi_U(\mathfrak{g}(x))$ is a uniformly bounded distance from the projection of $\pi_U(x)$ to $\pi_U(\gamma)$.

We will use $\mathfrak{g}$ as the map to prove γ is contracting. Gate maps satisfy condition (1) of Definition 4.2 by definition and condition (2) since they are coarsely Lipschitz. Hence it remains to prove that condition (3) of Lemma 4.3 holds.

Fix a point $x \in \mathcal{X}$ with $\mathsf{d}_{\mathcal{X}}(x, \gamma) \geqslant B_0$ and let $y \in \mathcal{X}$ be any point with $\mathsf{d}_{\mathcal{X}}(x, y) < B_1 \mathsf{d}_{\mathcal{X}}(x, \gamma)$ for constants B_0 and B_1 as determined below.

Since $\mathfrak{g}$ is a gate map and γ has D–bounded projections, for all $U \in \mathfrak{S} - \{S\}$ we have $\mathsf{d}_U(\mathfrak{g}(x), \mathfrak{g}(y)) < D$. Thus, by taking a threshold L for the distance formula (Theorem 2.5) larger than D, we have

$$\mathsf{d}_{\mathcal{X}}(\mathfrak{g}(x), \mathfrak{g}(y)) \asymp_{(K,C)} \mathsf{d}_S(\mathfrak{g}(x), \mathfrak{g}(y)),$$

for uniform constants K, C. Thus it suffices to prove that $\mathsf{d}_S(\mathfrak{g}(x), \mathfrak{g}(y))$ is bounded by some uniform constant B_2. We also choose L to be larger than the constants in Definition 2.1.(4).

By Definition 2.1.(1), the maps π_U are Lipschitz with a uniform constant. Taking B_0 sufficiently large, it follows that there exists $U \in \mathfrak{S}$ so that $\mathsf{d}_U(x, \mathfrak{g}(x)) > L$. By choosing B_1 to be sufficiently small, and applying the distance formula to the pairs (x, y) and $(x, \mathfrak{g}(x))$, the fact that the projections π_U are Lipschitz implies that the sum of the terms in the distance formula associated to $(x, \mathfrak{g}(x))$ is much greater than the sum of those associated to (x, y). Having chosen $B_1 < \frac{1}{2}$, we have $\sum \mathsf{d}_U(x, \mathfrak{g}(x)) > 2 \sum \mathsf{d}_U(x, y) > \sum(\mathsf{d}_U(x, y) + L)$. Thus, there exists $W \in \mathfrak{S}$ for which $\mathsf{d}_W(x, \mathfrak{g}(x)) > \mathsf{d}_W(x, y) + L$.

If $W = S$, then having $\mathsf{d}_S(x, \mathfrak{g}(x)) > \mathsf{d}_S(x, y) + L$ (where we enlarge L if necessary) would already show that the $\mathcal{C}S$–geodesic between x and y was disjoint from $\pi_S(\gamma)$ and then hyperbolicity of $\mathcal{C}S$ would yield a uniform bound on the $\mathsf{d}_S(\mathfrak{g}(x), \mathfrak{g}(y))$.

Otherwise, we may assume $W \neq S$. By the triangle inequality, we have $\mathsf{d}_W(y, \mathfrak{g}(x)) > L$. Further, since, as noted above, the $\mathcal{C}W$ projections between $\mathfrak{g}(x)$ and $\mathfrak{g}(y)$ are uniformly bounded, by choosing B_0 large enough and B_1 small enough, we also have $\mathsf{d}_W(y, \mathfrak{g}(y)) > L$.

By the bounded geodesic image axiom (Definition 2.1.(7)), any geodesic in $\mathcal{C}S$ either has bounded projection to $\mathcal{C}U$ or satisfies $\pi_S(\gamma) \cap \mathcal{N}_E(\rho^U_S) \neq \emptyset$ for any $U \in \mathfrak{S} - \{S\}$. For any geodesic from $\pi_S(x)$ to $\pi_S(\mathfrak{g}(x))$ (or from $\pi_S(y)$ to $\pi_S(\mathfrak{g}(y))$, the above argument implies that the first condition doesn't hold for W. Thus, in both cases, we know that any such geodesic must pass uniformly close to ρ^W_S. Hence the hyperbolicity of $\mathcal{C}S$ implies γ is contracting, and the first implication holds.

We prove the second implication by contradiction. By Theorem 3.7, we obtain a new structure $(\mathcal{X}, \mathfrak{T})$ which has unbounded products. For every $U \in \mathfrak{T} - \{S\}$ we have that both $\mathbf{F}_U$ and $\mathbf{E}_U$ are unbounded, hence every $U \in \mathfrak{T} - \{S\}$ yields a non-trivial product region $\mathbf{P}_U = \mathbf{E}_U \times \mathbf{F}_U$ which is uniformly quasi-isometrically embedded in $\mathcal{X}$.

Suppose γ is contracting but doesn't have D–bounded projections. Then we obtain a sequence $\{U_i\} \in \mathfrak{T} - \{S\}$ with $\mathsf{diam}(\pi_{\mathcal{C}U_i}(\gamma)) \to \infty$. Thus there is a sequence of pairs of points $x_i, y_i \in \gamma$, so that $d_{U_i}(x_i, y_i) \asymp K_i$, with $K_i \to \infty$. For each i, let q_i be a R–hierarchy path between x_i, y_i. By [BHS19, Proposition 5.17], there exists $\nu > 0$ depending only on R and $(\mathcal{X}, \mathfrak{S})$, such that

$$\mathsf{diam}_{U_i}(q_i \cap \mathcal{N}_\nu(\mathbf{P}_{U_i})) \asymp K_i.$$

Since γ is contracting, it is uniformly stable by Lemma 4.3. Since γ is uniformly stable and the q_i are uniform quasigeodesics, it follows that each q_i is contained in a uniform neighborhood of γ. Hence arbitrarily long segments of γ are uniformly close to the product regions $\mathbf{P}_{U_i}$. This contradicts the assumption that γ is contracting and completes the proof. □

5. Universal and largest acylindrical actions

The goal of this section is to show that for every hierarchically hyperbolic group $(G, \mathfrak{S})$ the poset $\mathcal{AH}(G)$ has a largest element. Recall that the action associated to such an element is necessarily a universal acylindrical action.

We prove the following stronger result which, in addition to providing new largest and universal acylindrical actions for cubulated groups, gives a single construction that recovers all previously known largest and universal acylindrical actions of finitely presented groups that are not relatively hyperbolic.

The following is Theorem A of the introduction:

Theorem 5.1. *Every hierarchically hyperbolic group admits a largest acylindrical action.*

Before giving the proof, we record the following result which gives a sufficient condition for an action to be largest. This result follows directly from the proof of [ABO19, Theorem 4.13]; we give a sketch of the argument here. Recall that an action $H \curvearrowright S$ is *elliptic* if H has bounded orbits.

Proposition 5.2 ([ABO19])**.** *Let G be a group, $\{H_1, \dots, H_n\}$ a finite collection of subgroups of G, and F be a finite subset of G such that $F \cup (\bigcup_{i=1}^n H_i)$ generates G. Assume that:*
(1) $\Gamma(G, F \cup (\bigcup_{i=1}^n H_i))$ is hyperbolic and the action of G on it is acylindrical.
(2) Each H_i is elliptic in every acylindrical action of G on a hyperbolic space.
Then $[F \cup (\bigcup_{i=1}^n H_i)]$ is the largest element in $\mathcal{AH}(G)$.

Proof. First notice that by assumption (1), $\Gamma(G, F \cup (\bigcup_{i=1}^n H_i)$ is an element of $\mathcal{AH}(G)$. Let $G \curvearrowright S$ be a cobounded acylindrical action of G on a hyperbolic space, S, and fix a basepoint $s \in S$. Then there exists a bounded subspace $B \subset S$ such that $S \subseteq \bigcup_{g \in G} g \cdot B$. By assumption (2), the orbit $H_i \cdot s$ is bounded for all $i = 1, \dots, n$. Since $|F| < \infty$, we know $\mathsf{diam}(F \cdot s) < \infty$ and thus

$$K = \max\{\mathsf{diam}(B), \mathsf{diam}(H_1 \cdot s), \dots, \mathsf{diam}(H_n \cdot s), \mathsf{diam}(F \cdot s)\}$$

is finite. Let $C = \{s' \in S \mid \mathsf{d}(s', s) \leqslant 3K\}$, and let

$$Z = \{g \in G \mid g \cdot C \cap C \neq \varnothing\}.$$

The standard Milnor-Schwartz Lemma argument shows that Z is an infinite generating set of G and there exists a G–equivariant quasi-isometry $S \to \Gamma(G, Z)$. It is clear that Z contains F, as well as H_i for all $i = 1, \dots, n$ and thus $[Z] \preceq [F \cup (\bigcup_{i=1}^n H_i]$. The result follows. □

Proof of Theorem 5.1. Let $(G, \mathfrak{S})$ be a hierarchically hyperbolic group with finite generating set F. By Corollary 3.8, there is a hierarchically hyperbolic group structure $(G, \mathfrak{T})$ with unbounded products. Recall that S is the $\sqsubseteq$–maximal element of $\mathfrak{T}$ with associated hyperbolic space $\mathcal{T}_S$. The action on $\mathcal{T}_S$ is acylindrical by [BHS17b, Theorem K].

Moreover, the action of G on $\mathcal{T}_S$ is cobounded, so let B be a fundamental domain for $G \curvearrowright \mathcal{T}_S$ and

$$\mathcal{U} = \{U \in \mathfrak{T} \mid \pi_S(\mathbf{F}_U) \subset B \text{ and } U \text{ is } \sqsubseteq\text{–maximal in } \mathfrak{T} - \{S\}\}.$$

Notice that $\mathcal{U}$ will contain exactly one representative from each G–orbit of domains, and so must be a finite set. Indeed, for a hierarchically hyperbolic group, this follows from the fact that the action of G on $\mathfrak{T}$ is cofinite.

Let $H_i \leqslant G$ be the stabilizer of $\mathbf{F}_{U_i}$ for each $U_i \in \mathcal{U}$. By a standard Milnor-Schwartz argument (see [ABO19] for details) there is a G–equivariant quasi-isometry between $\Gamma(G, F \cup (\bigcup_{i=1}^n H_i))$ and $\mathcal{T}_S$, where $n = |\mathcal{U}|$. Therefore condition (1) of Proposition 5.2 is satisfied.

By definition, each H_i sits inside a non-trivial direct product in G, the product region $\mathbf{P}_{U_i}$ associated to each $U_i \in \mathcal{U}$. It follows that H_i must be elliptic in every acylindrical action of G on a hyperbolic space (see Remark 2.18), satisfying condition (2).

Therefore, by Proposition 5.2, the action is largest. □

Remark 5.3. The proof of Theorem 5.1 can be extended to treat a number of groups which are hierarchically hyperbolic spaces, but not hierarchically hyperbolic groups. For example, it was shown in [BHS19, Theorem 10.1] that every fundamental group of a compact 3–manifold with no Nil or Sol in its prime decomposition admits a hierarchically hyperbolic space structure, which is constructed by first putting an HHS structure on each geometric piece in the prime decomposition. However, as explained in [BHS19, Remark 10.2] it is likely that such fundamental groups don't all admit hierarchically hyperbolic group structures. Nonetheless, the proof of the above theorem works in this case by replacing the use of the fact that the action of G on $\mathfrak{T}$ is cofinite, with the fact that for $\pi_1(M)$, the set $\mathcal{U}$ is precisely the set of $\sqsubseteq$–maximal domains in the hierarchically hyperbolic structure on each of the Seifert-fibered components of the prime decomposition of M, and so is finite.

Remark 5.4. There is an instructive direct proof of the universality of the above action, using the characterization of contracting quasigeodesics in Section 4, which we now give. We call an infinite order element *contracting* if its orbit is a contracting quasigeodesic in the Cayley graph. Now, let $g \in G$ be an infinite order element and consider the geodesic $\langle g \rangle$ in $\Gamma(G, F)$.

If $\langle g \rangle$ is contracting in $\Gamma(G, F)$, then by Theorem 4.4 all proper projections are bounded, and thus by the distance formula, g is loxodromic for the action on $\mathcal{T}_S$.

If $\langle g \rangle$ is not contracting in $\Gamma(G, F)$, then there exists some $U \in \mathfrak{T}$ such that $\pi_U(\langle g \rangle)$ is unbounded. Thus for any increasing sequence of constants (K_i) with $K_i \to \infty$, there are sequences of pairs of points $x_i, y_i \in \langle g \rangle$ such that $\mathsf{d}(x_i, y_i) \to \infty$ as $i \to \infty$ and $\mathsf{d}_U(x_i, y_i) \geqslant K_i$. For each i, let γ_i be an R–hierarchy path between x_i and y_i. By definition, γ_i is a uniform quasigeodesic. Then by [BHS19, Proposition 5.17], there exists $\nu > 0$ depending only on R and $(\mathcal{X}, \mathfrak{T})$ such that $\mathsf{diam}_U(\gamma_i \cap N_\nu(\mathbf{P}_U)) \geqslant K_i$. If g is a generalized loxodromic, then $\langle g \rangle$ is stable, by [Sis16], and so the subgeodesic $[x_i, y_i]$ stays within a uniform bounded distance of γ_i. Thus arbitrarily long subgeodesics of $\langle g \rangle$ stay within a uniformly bounded distance of a product region, $\mathbf{P}_U$. This contradicts $\langle g \rangle$ being Morse, and therefore g is not a generalized loxodromic element.

This remark directly implies that the action on $\mathcal{T}_S$ is a universal acylindrical action. (The universality of the action can also be proven using the classification of elements of $\mathrm{Aut}(\mathfrak{S})$ described in [DHS17].)

Another immediate consequence of the above remark is the following, which for hierarchically hyperbolic groups strengthens a result obtained by combining Osin [Osi16, Theorem 1.4.(L4)] and Sisto [Sis16, Theorem 1], which together prove that a generalized loxodromic element in an acylindrically hyperbolic group is quasi-geodesically stable.

Corollary 5.5. *Let $(G, \mathfrak{S})$ be a hierarchically hyperbolic group. An element $g \in G$ is generalized loxodromic if and only if g is contracting.*

The next result provides information about the partial ordering of acylindrical actions. Of the groups listed below, the largest and universal acylindrical action of the class of special CAT(0) cubical groups is new; the other cases were recently established to be largest in [ABO19].

Corollary 5.6. *The following groups admit acylindrical actions that are largest (and therefore universal):*

(1) Hyperbolic groups.
(2) Mapping class groups of orientable surfaces of finite type.
(3) Fundamental groups of compact three-manifolds with no Nil or Sol in their prime decomposition.
(4) Groups that act properly and cocompactly on a special CAT(0) cube complex, and more generally any cubical group which admits a factor system. This includes right angled-Artin groups, right-angled Coxeter groups, and many other examples as in [HS16].

Proof. With the exception of (3), by [BHS17b, BHS19, HS16] the above are all hierarchically hyperbolic groups and therefore have the bounded domain dichotomy. In case (3), where G is the fundamental group of a compact three-manifold with no Nil or Sol in its prime decomposition, then while G is not always a hierarchically hyperbolic group, it has a hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{S})$. To see this, we use the fact that there is a group G' which is quasi-isometric to G and has a hierarchically hyperbolic structure with all of the associated hyperbolic spaces infinite [BHS19, Theorem 10.1 & Remark 10.2]; thus by quasi-isometric invariance of hierarchically hyperbolic structures [BHS19, Proposition 1.10], G does as well. Since all of the associated hyperbolic spaces are infinite, $(\mathcal{X}, \mathfrak{S})$ has the bounded domain dichotomy, so the result follows. □

We give an explicit description of these actions for each hierarchically hyperbolic group in the corollary, in the sense that we describe the set $\mathfrak{W}$ of domains which are removed from the standard hierarchical structure of the group and whose associated hyperbolic space is infinite diameter. Recall that the space $\mathcal{T}_S$ is constructed from $\mathcal{X}$ by coning off all elements of $\mathfrak{T}$ which consists of those components of $\mathfrak{S}$ whose associated product regions have both factors with infinite diameter. Coning off all of $\mathfrak{T}$ yields a space which is is quasi-isometric to the space obtained by just coning off $\mathfrak{S} - \mathfrak{W}$.

(1) Hyperbolic groups have a canonical simplest hierarchically hyperbolic group structure given by taking $\mathfrak{S} = \{S\}$, where $\mathcal{C}S$ is the Cayley graph of the group with respect to a finite generating set. For this structure, $\mathfrak{W} = \varnothing$, and the action on the Cayley graph is clearly largest.
(2) For mapping class groups, the natural hierarchically hyperbolic group structure is $\mathfrak{S}$ is the set of homotopy classes of non-trivial non-peripheral (possibly disconnected) subsurfaces of the surface; the maximal element S is the surface itself, and the hyperbolic space $\mathcal{C}S$ is the curve complex. For this structure $\mathfrak{W} = \varnothing$. (Note that to form $\mathfrak{T}$ one must remove the nest-maximal collections of disjoint subsurfaces; the hyperbolic space associated to each of these, except S, has finite diameter). Additionally, we emphasize that although the new hyperbolic space $\mathcal{T}_S$ is not $\mathcal{C}S$, it is quasi-isometric to $\mathcal{C}S$, the action on which is known to be universal. Universality of this action was shown by Osin in [Osi16], and follows from results of Masur-Minsky and Bowditch [Bow08, MM99].
(3) If M is a compact 3–manifold with no Nil or Sol in its prime decomposition and $G = \pi_1 M$, then $\mathfrak{W}$ is exactly the set of vertex groups in the prime decomposition that are fundamental groups of hyperbolic 3–manifolds (each of which has exactly one domain in its hierarchically hyperbolic structure).
(4) If G is a group that acts properly and cocompactly on a special CAT(0) cube complex X, then by [BHS17b, Proposition B], X has a G–equivariant factor system. This factor system gives a hierarchically hyperbolic group structure in which $\mathfrak{S}$ is the closure under projection of the set of hyperplanes along with a maximal element S, where $\mathcal{C}S$ is the contact graph as defined in [Hag14]. In this structure, $\mathfrak{W}$ is the set of indices whose stabilizer in G contains a power of a rank one element.

    In the particular case of right-angled Artin groups, no power of a rank one element will stabilize a hyperplane, so $\mathfrak{W} = \varnothing$. In this case, the contact graph $\mathcal{C}S$ is quasi-isometric to the extension graph defined by [KK14]. That the action on the extension graph is a universal acylindrical action follows from the work of [KK14] and the centralizer theorem for right-angled Artin groups. This action is also shown to be largest in [ABO19].

We give a concrete example of the situation in the case of a right-angled Coxeter group.

Example 5.7. Let G be the right-angled Coxeter group whose defining graph is a pentagon. Then $G = \langle a, b, c, d, e \mid [a, b], [b, c], [c, d], [d, e], [a, e], a^2, b^2, c^2, d^2, e^2 \rangle$, and the Cayley graph of G is the tiling of the hyperbolic plane by pentagons. We consider the dual square complex to this tiling. To form the contact graph $\mathcal{C}S$, we start with the square complex and cone off each hyperplane carrier, which is equivalent to coning off the hyperplane stabilizers in the Cayley graph. The result is a quasi-tree. Thus a fundamental domain for the hierarchically hyperbolic group structure of G is $\{U_a, U_b, U_c, U_d, U_e, S\}$ where U_v is associated to the stabilizer of the hyperplane labeled by v and S is associated to the contact graph described above.

Consider the hyperplane J_b that is labeled by b. Then the stabilizer of J_b is subgroup generated by the star of the vertex b, that is $\langle a, b, c\rangle$. This subgroup contains the infinite order element ac. As G is a hyperbolic group, all infinite order elements are generalized loxodromic, but ac is not loxodromic for the action on the contact graph since its axis lies in a hyperplane stabilizer that has been coned-off. Thus the action on the contact graph is not universal.

Let $U_b \in \mathfrak{S}$ be the element associated to $\mathrm{Stab}(J_b)$. Then $\mathrm{Stab}(J_b) = \langle a, b, c \mid [a, b], [b, c]\rangle \simeq D_\infty \times \mathbb{Z}/2\mathbb{Z} \simeq \mathbf{F}_{U_b} \times \mathbf{E}_{U_b}$ is a product region, and the maximal orthogonal component $\mathbf{E}_{U_b}$ is bounded. Thus $U_b \in \mathfrak{W}$, as is U_v for each vertex v of the defining graph. The contact graph associated to $(\mathbf{F}_{U_b}, \mathfrak{S}_{U_b})$ is a line, and the element ac is loxodromic for the action on this space.

Note that once $\mathfrak{W}$ has been removed from $\mathfrak{S}$, the resulting hierarchically hyperbolic structure is $(G, \{S\})$, the canonical hierarchically hyperbolic structure for a hyperbolic group, in which $\mathcal{C}S = \Gamma(G, \{a, b, c, d, e\})$.

6. Characterizing stability

In this section, we will give several characterizations of stability which hold in any hierarchically hyperbolic group. In fact, we will characterize stable embeddings of geodesic metric spaces into hierarchically hyperbolic spaces with unbounded products. One consequence of this will be a description of points in the Morse boundary of a proper geodesic hierarchically hyperbolic space with unbounded products as the subset of the hierarchically hyperbolic boundary consisting of points with bounded projections.

6.1. Stability.

While it is well-known that contracting implies stability [Beh06, DMS10, MM99], the converse is not true in general. Nonetheless, in several important classes of spaces the converse holds, including in hyperbolic spaces, CAT(0) spaces, the mapping class group, and Teichmüller space [Sul14, Beh06, DT15, Min96]. We record the following corollary of Theorem 4.4 which gives a relationship between stability and contracting subsets that holds in a fairly general context.

Corollary 6.1. *Suppose that $(\mathcal{X}, \mathfrak{S})$ has unbounded products, $\mathcal{Y}$ is a hyperbolic metric space, and $i\colon \mathcal{Y} \to \mathcal{X}$ is a (K, C)–quasi-isometric embedding. Then $i(\mathcal{Y})$ is N–stable if and only if $i(\mathcal{Y})$ is D–contracting, where N and D determine each other.*

Proof. First assume that $i(\mathcal{Y})$ is D–contracting. Since $i\colon \mathcal{Y} \to \mathcal{X}$ is a (K, C)–quasi-isometric embedding, to show that $i(\mathcal{Y})$ is N–stable for some gauge $N = N(D)$, we need only show that the (quasigeodesic) image $i(\gamma)$ of every geodesic γ in $\mathcal{Y}$ is $N(K, C)$–stable. Since $i(\mathcal{Y})$ is D–contracting and $i(\mathcal{Y})$ is hyperbolic, $i(\gamma)$ is D'–contracting for some D' depending only on D, K, C, and the hyperbolicity constant of $\mathcal{Y}$. Lemma 4.3 shows that $i(\gamma)$ is therefore N–stable, with N depending only on D, as desired. (Note that the assumption on unbounded products is not necessary for this implication.)

For the other direction, the fact that $\mathcal{X}$ has unbounded products implies that $i(\mathcal{Y})$ has bounded projections, since otherwise one could find large segments of quasigeodesics contained inside product regions with unbounded factors, contradicting stability. The result now follows from Theorem 4.4. □

The following provides a general characterization of stability in HHSs, a special case of which is Theorem B.

Corollary 6.2. *Let $i\colon \mathcal{Y} \to \mathcal{X}$ be a quasi-isometric embedding from a metric space into a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ with unbounded products. The following are equivalent:*

(1) i is a stable embedding;
(2) $i(\mathcal{Y})$ has uniformly bounded projections;
(3) $\pi_S \circ i\colon \mathcal{Y} \to \mathcal{C}S$ is a quasi-isometric embedding.

Proof. That item (2) implies (3) follows from the distance formula and the assumption that i is a quasi-isometric embedding.

The hypothesis of item (1) implies that $\mathcal{Y}$ is hyperbolic. Moreover, since (2) implies (3), the hypothesis of (2) also implies that $\mathcal{Y}$ is hyperbolic. Thus items (1) and (2) are equivalent via Corollary 6.1 and Theorem 4.4.

We now prove that (3) implies (2). Suppose for a contradiction that for any integer N there exists $U \in \mathfrak{S} - \{S\}$ and $x, y \in i(\mathcal{Y})$ satisfying $d_U(x, y) > N$. Now, we consider a hierarchy path γ between x and y. Applying the bound geodesic image axiom (Definition 2.1.(7)) to the associated $\mathcal{C}S$–geodesic between $\pi_S \circ i(x)$ and $\pi_S \circ i(y)$ it follows that this $\mathcal{C}S$–geodesic has non-trivial intersection with the radius E ball about the point ρ^U_S. Indeed, this yields that there exist points x', y' on the geodesic which are both distance at most E from ρ^U_S; by [BHS19, Lemma 5.17] we can assume that x and y were chosen so that x' and y' also satisfy $d_S(x, x') < E$ and $d_S(y, y') < E$. Thus, we have that $d_S(x, y) < 4E$. The hypothesis in (3) implies that there is a uniform bound on $d_{\mathcal{Y}}(x, y)$. The distance formula then implies a uniform bound on $d_W(x, y)$ for any $W \in \mathfrak{S}$, contradicting the fact that we chose $d_U(x, y)$ to be large. □

6.2. **The Morse boundary.** In the rest of this section, we turn to studying the Morse boundary and use this to give a bound on the stable asymptotic dimension of a hierarchically hyperbolic space. We begin by describing two notions of boundary.

In [DHS17], Durham, Hagen, and Sisto introduced a boundary for any hierarchically hyperbolic space. We collect the relevant properties we need in the following theorem:

Theorem 6.3 (Theorem 3.4 and Proposition 5.8 in [DHS17])**.** *If $(\mathcal{X}, \mathfrak{S})$ is a proper hierarchically hyperbolic space, then there exists a topological space $\partial\mathcal{X}$ such that $\partial\mathcal{X} \cup \mathcal{X} = \overline{\mathcal{X}}$ compactifies $\mathcal{X}$, and the action of $\operatorname{Aut}(\mathcal{X}, \mathfrak{S})$ on $\mathcal{X}$ extends continuously to an action on $\overline{\mathcal{X}}$.*

Moreover, if $\mathcal{Y}$ is a hierarchically quasiconvex subspace of $\mathcal{X}$, then, with respect to the induced hierarchically hyperbolic structure on $\mathcal{Y}$, the limit set of $\Lambda\mathcal{Y}$ of $\mathcal{Y}$ in $\partial\mathcal{X}$ is homeomorphic to $\partial\mathcal{Y}$ and the inclusion map $i\colon \mathcal{Y} \to \mathcal{X}$ extends continuously an embedding $\partial i\colon \partial\mathcal{Y} \to \partial\mathcal{X}$.

Building on ideas in [CS15], Cordes introduced the *Morse boundary* of a proper geodesic metric space [Cor17], which was then refined further by Cordes–Hume in [CH17]. The *Morse boundary* is a stratified boundary which encodes the asymptotic classes of Morse geodesic rays based at a common point. Importantly, it is a quasi-isometry invariant and generalizes the Gromov boundary of a hyperbolic space [Cor17].

We briefly discuss the construction of the Morse boundary and refer the reader to [Cor17, CH17] for details.

Consider a a proper geodesic metric space X with a basepoint $e \in X$. Given a stability gauge $N\colon \mathbb{R}^2_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$, define a subset $X^{(N)}_e \subset X$ to be the collection of points $y \in X$ such that e and y can be connected by an N–stable geodesic in X. Each such $X^{(N)}_e$ is δ_N–hyperbolic for some $\delta_N > 0$ depending on N and X [CH17, Proposition 3.2]; here, we use the Gromov product definition of hyperbolicity, as $X^{(N)}_e$ need not be connected. Moreover, any stable subset of X embeds in $X^{(N)}_e$ for some N [CH17, Theorem A.V].

The set of stability gauges admits a partial order: $N_1 < N_2$ if and only if $N_1(K, C) < N_2(K, C)$ for all constants K, C. In particular, if $N_1 < N_2$, then $X^{(N_1)}_e \subset X^{(N_2)}_e$.

Since each $X_e^{(N)}$ is Gromov hyperbolic, each admits a Gromov boundary $\partial X_e^{(N)}$. Take the direct limit with respect to this partial order to obtain a topological space $\partial_s X$ called the *Morse boundary* of X.

We fix $(\mathcal{X}, \mathfrak{S})$, a hierarchically hyperbolic structure with unbounded products.

Definition 6.4. We say $\lambda \in \partial\mathcal{X}$ has *bounded projections* if for any $e \in \mathcal{X}$, there exists $D > 0$ such that any R–hierarchy path $[e, \lambda]$ has D–bounded projections. Let $\partial_c\mathcal{X}$ denote the set of points $\lambda \in \partial\mathcal{X}$ with bounded projections.

The boundary $\partial\mathcal{X}$ contains $\partial\mathcal{C}U$ for each $U \in \mathfrak{S}$, by construction. The next lemma shows that the boundary points with bounded projections are contained in $\partial\mathcal{C}S$, as a subset of $\partial\mathcal{X}$, where S is the $\sqsubseteq$–maximal element. In general, the set of boundary points with bounded projections may be a very small subset of $\partial\mathcal{C}S$. For instance, in the boundary of the Teichmüller metric, these points are a proper subset of the uniquely ergodic ending laminations and have measure zero with respect to any hitting measure of a random walk on the mapping class group.

Lemma 6.5. *The inclusion $\partial_c\mathcal{X} \subset \partial\mathcal{C}S$ holds for any $(\mathcal{X}, \mathfrak{S})$ with unbounded products where S is the $\sqsubseteq$–maximal element of $\mathfrak{S}$. Moreover, if $\mathcal{X}$ is also proper, then for any $D > 0$ there exists $D' > 0$ depending only on D and $(\mathcal{X}, \mathfrak{S})$ such that if $(x_n) \subset \mathcal{X}$ is a sequence with $x_n \to \lambda \in \partial\mathcal{X}$ such that $[e, x_n]$ has D–bounded projections for some $e \in \mathcal{X}$ and each n, then $[e, \lambda]$ has D'–bounded projections.*

Proof. Let $\lambda \in \partial_c\mathcal{X}$. If $[e, \lambda]$ is an R–hierarchy path, then $[e, \lambda]$ has an infinite diameter projection to some $\mathcal{C}U$, see, e.g., [DHS17, Lemma 3.3]. As λ has bounded projections, we must have $U = S$. Since $\pi_S([e, \lambda]) \subset \mathcal{C}S$ is a quasigeodesic ray, the first statement follows.

Now suppose that $\mathcal{X}$ is also proper. For each n, let $\gamma_n = [e, x_n]$ be any R–hierarchy path between e and x_n in $\mathcal{X}$. The Arzela-Ascoli theorem implies that after passing to a subsequence, γ_n converges uniformly on compact sets to some R'–hierarchy path γ with R' depending only on R and $(\mathcal{X}, \mathfrak{S})$. Hence γ has D'–bounded projections for some D' depending only on D and $(\mathcal{X}, \mathfrak{S})$. Moreover, since $x_n \to \lambda$ in $\mathcal{C}S$, it follows that $\pi_S(\gamma)$ is asymptotic to λ in $\mathcal{C}S$.

If $[e, \lambda]$ is any other R'–hierarchy path, it follows from uniform hyperbolicity of the $\mathcal{C}U$ and the definition of hierarchy paths that $d_U^{Haus}(\gamma, [e, \lambda])$ is uniformly bounded for all $U \in \mathfrak{S}$. Since γ has D'–bounded projections, the distance formula implies that $[e, \lambda]$ has D''–bounded projections for some D'' depending only on D and $(\mathcal{X}, \mathfrak{S})$, as required. □

6.3. **Bounds on stable asymptotic dimension.** The asymptotic dimension of a metric space is a coarse notion of topological dimension which is invariant under quasi-isometry. Introduced by Cordes–Hume [CH17], the stable asymptotic dimension of a metric space X is the maximal asymptotic dimension a stable subspace of X.

The stable asymptotic dimension of a metric space X is always bounded above by its asymptotic dimension. Behrstock, Hagen, and Sisto [BHS17a] proved that all proper hierarchically hyperbolic spaces have finite asymptotic dimension (and thus have finite stable asymptotic dimension, as well). The bounds on asymptotic dimension obtained in [BHS17a] are functions of the asymptotic dimension of the top level curve graph.

In the following theorem, we prove that a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ has finite stable asymptotic dimension under the assumption that $\operatorname{asdim}(\mathcal{C}S) < \infty$, where $\mathcal{C}S$ is the hyperbolic space associated to the $\sqsubseteq$–maximal domain S in $\mathfrak{S}$.

Recall that asymptotic dimension is monotonic under taking subsets. Thus, if $\mathcal{X}$ is assumed to be proper, so that $\operatorname{asdim}(\mathcal{C}S) < \infty$, then $\mathcal{X}$ (and therefore its stable subsets) have finite asymptotic dimension by [BHS17a]. Here, using some geometry of stable subsets we obtain a sharper bound on $\operatorname{asdim}_s(\mathcal{X})$ than $\operatorname{asdim}(\mathcal{X})$.

Theorem 6.6. *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space with unbounded products such that $\mathcal{C}S$ has finite asymptotic dimension, where S is the $\sqsubseteq$–maximal element of $\mathfrak{S}$. Then* $\operatorname{asdim}_s(\mathcal{X}) \leqslant \operatorname{asdim}(\mathcal{C}S)$. *Moreover, if $\mathcal{X}$ is also proper and geodesic, then there exists a continuous bijection $\widehat{i}\colon \partial_s\mathcal{X} \to \partial_c\mathcal{X}$.*

Proof. By [CH17, Lemma 3.6], for any stability gauge N there exists N' such that $\mathcal{X}_e^{(N)}$ is N'–stable. Hence, there exists $D' > 0$ depending only on N' and $(\mathcal{X}, \mathfrak{S})$ such that $\mathcal{X}_e^{(N)}$ has D'–bounded projections. By Corollary 6.2, it follows that the projection $\pi_S\colon \mathcal{X}_e^{(N)} \to \mathcal{C}S$ is a quasi-isometric embedding with constants depending only on D' and $(\mathcal{X}, \mathfrak{S})$. Since every stable subset of $\mathcal{X}$ embeds into some $\mathcal{X}_e^{(N)}$ [CH17, Theorem A.V], the first conclusion then follows from the definition of stable asymptotic dimension.

Now suppose that $\mathcal{X}$ is proper.

Since each $\mathcal{X}_e^{(N)}$ is stable in $\mathcal{X}$, these sets have bounded projections by Corollary 6.2; from this it follows that $\mathcal{X}_e^{(N)}$ is hierarchically quasiconvex for each N. Hence by [DHS17, Proposition 5.8], the canonical embedding $i^{(N)}\colon \mathcal{X}_e^{(N)} \hookrightarrow \mathcal{X}$ extends to an embedding $\widehat{i}^{(N)}\colon \partial\mathcal{X}_e^{(N)} \hookrightarrow \partial\mathcal{X}$.

By Corollary 6.2 and Lemma 6.5, we have $\widehat{i}^{(N)}\left(\partial\mathcal{X}_e^{(N)}\right) \subset \partial_c\mathcal{X} \subset \partial\mathcal{C}S$. Let $\widehat{i} : \partial_s\mathcal{X} \to \partial_c\mathcal{X}$ be the direct limit of the $\widehat{i}^{(N)}$. Since it is injective on each stratum, $\widehat{i}$ is injective.

To prove surjectivity, let $\lambda \in \partial_c\mathcal{X}$. Let $e \in \mathcal{X}$ and fix a hierarchy path $[e, \lambda]$. Since $\lambda \in \partial_c\mathcal{X}$, $[e, \lambda]$ has D–bounded projections for some $D > 0$. Let $x_n \in [e, \lambda]$ be such that $x_n \to \lambda$ in $\overline{\mathcal{X}}$. If $[e, x_n]$ is a sequence of geodesics between e and x_n, then, by properness, the Arzela–Ascoli theorem, and passing to a subsequence if necessary, there exists a geodesic ray $\gamma\colon [0, \infty) \to \mathcal{X}$ with $\gamma(0) = e$ such that $[e, x_n]$ converges on compact sets to γ. Since each $[e, x_n]$ has D–bounded projections, it follows that γ has D'–bounded projections for some D' depending only on D and $(\mathcal{X}, \mathfrak{S})$. Moreover, by hyperbolicity of $\mathcal{C}S$ and the construction of γ we have that $d_{\mathcal{C}S}^{Haus}(\pi_S(\gamma), [e, \lambda])$ is uniformly bounded and thus, by the distance formula, so is $d_{\mathcal{X}}^{Haus}(\gamma, [e, \lambda])$. Since $[(x_n)] = [\gamma]$ by construction, it follows that $\widehat{i}(\gamma) = \lambda$, as required.

Continuity of $\widehat{i}^{(N)}$ for each N follows from [DHS17, Proposition 5.8], as above. This and the definition of the direct limit topology implies continuity of $\widehat{i}$. $\square$

The following corollary is immediate:

Corollary 6.7. *If G is a hierarchically hyperbolic group, then G has finite stable asymptotic dimension.*

6.4. **Random subgroups.** Let G be any countable group and μ a probability measure on G whose support generates a non-elementary semigroup. A *k–generated random subgroup of G*, denoted $\Gamma(n)$ is defined to be the subgroup $\langle w_n^1, w_n^2, \ldots, w_n^k \rangle \subset G$ generated by the n^{th} step of k independent random walks on G, where $k \in \mathbb{N}$. For other recent results on the geometry of random subgroups of acylindrically hyperbolic groups, see [MS19].

Following Taylor-Tiozzo [TT16], we say a k–generated random subgroup $\Gamma(n)$ of G has a property P if

$$\mathbb{P}[\Gamma(n) \quad \text{has} \quad P] \to 1 \quad \text{as} \quad n \to \infty.$$

Theorem 6.8. *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS for which the $\sqsubseteq$–maximal element, S, has $\mathcal{C}S$ infinite diameter, and consider $G < \operatorname{Aut}(\mathcal{X}, \mathfrak{S})$ which acts properly and cocompactly on $\mathcal{X}$ via the orbit map. Then any k–generated random subgroup of G stably embeds in $\mathcal{X}$.*

Proof. By [BHS17b, Theorem K], G acts acylindrically on $\mathcal{C}S$. Let $\Gamma(n)$ be generated by k independent random walks as above. Now, [TT16, Theorem 1.2] implies that $\Gamma(n)$ a.a.s. quasi-isometrically embeds into $\mathcal{C}S$, and hence $\Gamma(n)$ is hyperbolic. Moreover, the distance formula implies that $\Gamma(n)$ is undistorted in G and any orbit of $\Gamma(n)$ in $\mathcal{X}$ has bounded projections by the distance formula. By Theorem 4.4, having bounded projections implies contracting; thus any orbit of $\Gamma(n)$ in $\mathcal{X}$ is a.a.s. contracting, which gives the conclusion by Corollary 6.1. (Note that the directions of Theorem 4.4 and Corollary 6.1 used here do not require that $(\mathcal{X}, \mathfrak{S})$ has unbounded products.) □

In particular, one consequence is a new proof of the following result of Maher–Sisto. This result follows from the above, together with Rank Rigidity for HHG ([DHS17, Theorem 9.14]) which implies that a hierarchically hyperbolic group which is not a direct product of two infinite groups has $\mathcal{CS}$ infinite diameter.

Corollary 6.9 (Maher–Sisto; [MS19])**.** *If G is a hierarchically hyperbolic group which is not the direct product of two infinite groups, then any k–generated random subgroup of G is stable.*

7. Clean containers

The clean container property is a condition related to the orthogonality axiom. In Proposition 7.2 this property is shown to hold for many groups, though it does not hold for all groups. Unlike earlier versions of this paper, this condition is no longer needed to prove the main theorems of the earlier sections. However, we keep the content in this paper since this property has found independent interest and is used elsewhere.

Definition 7.1 (Clean containers)**.** In a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ for each $T \in \mathfrak{S}$ and each $U \in \mathfrak{S}_T$ with $\{V \in \mathfrak{S}_T \mid V \perp U\} \neq \varnothing$ the orthogonality axiom provides a container. If, for each U, such a container can be chosen to be orthogonal to U, then we say that $(\mathcal{X}, \mathfrak{S})$ has *clean containers.*

We first describe some interesting examples with clean containers. Then we show that this property is preserved under some combination theorems for hierarchically hyperbolic spaces. We refer the reader to [BHS19, Sections 8 & 9] and [BHS17a, Section 6] for details on the structure in the new spaces.

Proposition 7.2. *The following spaces admit hierarchically hyperbolic structures with clean containers.*

1. *Hyperbolic groups*
2. *Mapping class groups of orientable surfaces of finite type*
3. *Special cubical groups, and more generally, any cubical group which admits a factor system.*
4. *$\pi_1(M)$, for M a compact 3–manifold with no Nil or Sol in its prime decomposition.*

Proof. Hierarchically hyperbolic structures for these spaces were constructed in [BHS17b] and [BHS19].

1. The statement is immediate for hyperbolic groups, as they all admit hierarchically hyperbolic structure with no orthogonality, and thus the container axiom is vacuous.
2. For mapping class groups, in the standard structure, a container for domains orthogonal to a given subsurface U is the complementary subsurface, which is orthogonal to U.
3. The statement follows immediately from [BHS17b, Proposition B] and [HS16, Corollary 3.4].
4. Given a geometric 3–manifold M of the above form, $\pi_1(M)$ is quasi-isometric to a (possibly degenerate) product of hyperbolic spaces, and so has clean containers by Proposition 7.3. Given an irreducible non-geometric graph manifold M, the hierarchically hyperbolic structure comes from considering $\pi_1(M)$ as a tree of hierarchically hyperbolic spaces with clean containers and hence has clean containers by Proposition 7.5. Finally, the general case of a non-geometric 3–manifold M follows

immediately from Proposition 7.4 and the fact that $\pi_1(M)$ is hyperbolic relative to its maximal graph manifold subgroups.

□

Proposition 7.3. *The product of two hierarchically hyperbolic spaces which both have clean containers has clean containers.*

Proof. Let $(\mathcal{X}_0, \mathfrak{S}_0)$ and $(\mathcal{X}_1, \mathfrak{S}_1)$ be hierarchically hyperbolic spaces with clean containers. In the hierarchically hyperbolic structure $(\mathcal{X}_0 \times \mathcal{X}_1, \mathfrak{S})$ given by [BHS19, Theorem 8.27] there are two types of containers, those that come from one of the original structures and those that do not. Containers of the first type are clean, as both original structures have clean containers.

The second type of domain consists of new domains obtained as follows. Given a domain $U \in \mathfrak{S}_i$, a new domain V_U is defined with the property that it contains under nesting any domain in $\mathfrak{S}_i$ which is orthogonal to U and also any domain in $\mathfrak{S}_{i+1}$. Thus, by construction V_U is a container for everything orthogonal to U. As $V_U \perp U$, the result follows. □

Proposition 7.4. *If G is hyperbolic relative to a collection of hierarchically hyperbolic spaces which all have clean containers, then G is a hierarchically hyperbolic space with clean containers.*

Proof. That G is a hierarchically hyperbolic space follows from [BHS19, Theorem 9.1]. In the hierarchically hyperbolic structure on G, no new orthogonality relations are introduced, and thus all containers are containers in the hierarchically hyperbolic structure of one of the peripheral subgroups. As each of these structures have clean containers, it follows that G does, as well. □

The following example relies on the combination theorem [BHS19, Theorem 8.6]. We provide this as another example of hierarchically hyperbolic spaces with clean containers, but since we don't rely on this elsewhere in the paper, we refer to that reference for the relevant definitions. Nonetheless, we include a full proof for the expert, since it is short. (We note that after this paper was circulated, Berlai and Robbio proved a combination theorem under weaker conditions than [BHS19, Theorem 8.6] and, in the process, also proved that if all the vertex spaces have clean containers, then so does the combined space, see [BR, Theorem A].)

Proposition 7.5. *Let $\mathcal{T}$ be a tree of hierarchically hyperbolic spaces satisfying the hypotheses of* [BHS19, Theorem 8.6]*, so that $X(\mathcal{T})$ is hierarchically hyperbolic. If for each $v \in T$, the hierarchically hyperbolic space $(\mathcal{X}_v, \mathfrak{S}_v)$ has clean containers, then so does $X(\mathcal{T})$.*

Proof. This follows immediately from the proof of [BHS19, Theorem 8.6] and the fact that edge-hieromorphisms are full and preserve orthogonality. In the notation from that result, we note that, if $\mathfrak{S}_v$ has clean containers for each $v \in T$, then the domain $A_v \in \mathfrak{S}_v$ described in the proof also has the property that $A_v \perp U_v$. Therefore, as edge-hieromorphisms are full and preserve orthogonality, $[A_v] \perp [U]$. □

The following uses the notion of *hierarchically hyperbolically embedded subgroups* introduced in [BHS17a]; see also [DGO17] for the related notion of hyperbolically embedded subgroups.

Proposition 7.6. *Let $(G, \mathfrak{S})$ be a hierarchically hyperbolic group with clean containers, and let H be a hierarchically hyperbolically embedded subgroup of $(G, \mathfrak{S})$. Then there exists a finite set $F \subset H - \{1\}$ such that for all $N \lhd H$ with $F \cap N = \emptyset$ and H/N is hyperbolic, the group $G/\hat{N}$, obtained by quotienting by the normal closure, is a hierarchically hyperbolic group with clean containers.*

Proof. Recall that in the hierarchically hyperbolic structure $(G/\hat{N}, \mathfrak{S}_N)$ obtained in [BHS17a, Theorem 6.2] (and in the notation used there), two domains $\mathbf{U}, \mathbf{V} \in \mathfrak{S}_{\mathbf{N}}$ satisfy $\mathbf{U} \sqsubseteq \mathbf{V}$ (respectively $\mathbf{U} \perp \mathbf{V}$) if there exists a linked pair $\{U, V\}$ with $U \in \mathbf{U}$ and $V \in \mathbf{V}$ such that $U \sqsubseteq V$ (respectively $U \perp V$). Let $\mathbf{T} \in \mathfrak{S}_{\mathbf{N}}$ and $\mathbf{U} \in (\mathfrak{S}_{\mathbf{N}})_{\mathbf{T}}$ with $\mathcal{V} = \{\mathbf{V} \in \mathfrak{S}_{\mathbf{T}} \mid \mathbf{V} \perp \mathbf{U}\} \neq \emptyset$. To prove the container axiom, we consider domains $T, U, V \in \mathfrak{S}$ such that $T \in \mathbf{T}$, $U \in \mathbf{U}$ and $V \in \mathbf{V}$ for all $\mathbf{V} \in \mathcal{V}$, and such that any pair is a linked pair. Then the orthogonality axiom for $(G, \mathfrak{S})$ provides a domain W such that $W \sqsupseteq V$ and $W \sqsubseteq T$. As $(G, \mathfrak{S})$ has clean containers, we also have that $W \perp U$. This implies that ρ^U_S and ρ^W_S are coarsely equal by [DHS17, Lemma 1.5], and so $\{U, W\}$ is a linked pair. Therefore, $\mathbf{W} \perp \mathbf{U}$. □

Appendix A. Almost HHSs are HHSs.

By Daniel Berlyne and Jacob Russell

The main result in this appendix is that every almost HHS structure can be promoted to an HHS structure. Recall that, as introduced in Section 3.2, an almost HHS is a space which satisfies all the axioms of an HHS except for the orthogonality axiom, which is instead replaced by a weaker axiom without a container requirement. In Theorem A.1, we show that an almost HHS structure can be made into an actual HHS structure by adding appropriately chosen "dummy domains" to serve as the containers. This result provides a useful method for producing an HHS structure while only needing to verify the weaker axioms of an almost HHS. This method is used in the main text in the proof of Theorem 3.7, where it is shown that every hierarchically hyperbolic space with the bounded domain dichotomy admits an HHS structure with unbounded products.

Theorem A.1. *Let $(\mathcal{X}, \mathfrak{S})$ be an almost HHS. There exists an HHS structure $\mathfrak{R}$ for $\mathcal{X}$ so that $\mathfrak{S} \subseteq \mathfrak{R}$, and if $W \in \mathfrak{R} - \mathfrak{S}$ then the associated hyperbolic space for W is a single point.*

To prove Theorem A.1, we will need to collect three additional tools about almost HHSs. Each of these tools was proved in the setting of hierarchically hyperbolic spaces, but they continue to hold in the almost HHS setting. Indeed, the only use of the containers in their proofs is [BHS19, Lemma 2.1], which proves that the cardinality of any collection of pairwise orthogonal domains is uniformly bounded by the complexity of the HHS.

The first tool says the relative projections of orthogonal domains coarsely coincide. Note, ρ^W_Q and ρ^V_Q are both defined when $W \pitchfork Q$ or $W \subsetneq Q$ and $V \pitchfork Q$ or $V \subsetneq Q$.

Lemma A.2 ([DHS17, Lemma 1.5])**.** *Let $(\mathcal{X}, \mathfrak{S})$ be an almost HHS. If $W, V \in \mathfrak{S}$ with $W \perp V$, and $Q \in \mathfrak{S}$ with ρ^W_Q and ρ^V_Q both defined, then $\mathsf{d}_Q(\rho^W_Q, \rho^V_Q) \leqslant 2\kappa_0$ where κ_0 is the constant from the consistency axiom of $\mathfrak{S}$.*

The second tool we will need is the realization theorem for almost HHSs. The realization theorem characterizes which tuples in the product $\prod_{V \in \mathfrak{S}} \mathcal{C}V$ are coarsely the image of a point in $\mathcal{X}$. Essentially, it says if a tuple $(b_V) \in \prod_{V \in \mathfrak{S}} \mathcal{C}V$ satisfies the consistency inequalities of an almost HHS (see Definition 2.6), then there exists a point $x \in \mathcal{X}$ such that $\pi_V(x)$ is uniformly close to b_V for each $V \in \mathfrak{S}$.

Theorem A.3 (The realization of consistent tuples, [BHS19, Theorem 3.1])**.** *Let $(\mathcal{X}, \mathfrak{S})$ be an almost HHS. There exists a function $\tau \colon [0, \infty) \to [0, \infty)$ so that if $(b_V)_{V \in \mathfrak{S}}$ is a κ–consistent tuple, then there exists $x \in \mathcal{X}$ so that $\mathsf{d}_V(x, b_V) \leqslant \tau(\kappa)$ for all $V \in \mathfrak{S}$.*

The last result we need is that the relative projections of an almost HHS also satisfy the inequalities in the consistency axiom.

Lemma A.4 (ρ–consistency, [BHS19, Proposition 1.8])**.** *Let $\mathfrak{S}$ be an almost HHS structure for $\mathcal{X}$ and $V, W, Q \in \mathfrak{S}$. Suppose $W \pitchfork Q$ or $W \subsetneq Q$ and $W \pitchfork V$ or $W \subsetneq V$. Then we have the following, where κ_0 is the constant from the consistency axiom of $(\mathcal{X}, \mathfrak{S})$.*

(1) *If* $Q \pitchfork V$, *then* $\min\{\mathsf{d}_Q(\rho^W_Q, \rho^V_Q), \mathsf{d}_V(\rho^Q_V, \rho^W_V)\} \leqslant 2\kappa_0$.
(2) *If* $Q \sqsubseteq V$, *then* $\min\{\mathsf{d}_V(\rho^Q_V, \rho^W_V), \mathsf{diam}(\rho^W_Q \cup \rho^V_Q(\rho^W_V))\} \leqslant 2\kappa_0$.

We are now ready to prove that every almost HHS is an HHS (Theorem A.1). If $(\mathcal{X}, \mathfrak{S})$ is an almost HHS, then the only HHS axiom that is not satisfied is the container requirement of the orthogonality axiom. The most obvious way to address this is to add an extra element to $\mathfrak{S}$ every time we need a container. That is, if $V, W \in \mathfrak{S}$ with $V \sqsubseteq W$ and there exists some $Q \sqsubseteq W$ with $Q \perp V$, then we add a domain D^V_W to serve as the container for V in W, i.e., every Q nested into W and orthogonal to V will be nested into D^V_W. However, this approach is perilous as once a domain Q is nested into D^V_W, we may now need a container for Q in D^V_W! To avoid this, we add domains $D^{\mathcal{V}}_W$ where $\mathcal{V}$ is a pairwise orthogonal set of domains nested into W; that is, $D^{\mathcal{V}}_W$ contains all domains Q that are nested into W and orthogonal to all $V \in \mathcal{V}$. This allows for all the needed containers to be added at once, avoiding an iterative process.

Proof of Theorem A.1. Let $(\mathcal{X}, \mathfrak{S})$ be an almost HHS and let $E \geqslant 0$ be the maximum of all the constants in $\mathfrak{S}$. Let $\mathcal{V}$ denote a non-empty set of pairwise orthogonal elements of $\mathfrak{S}$ and let $W \in \mathfrak{S}$. We say the pair $(W, \mathcal{V})$ is a *container pair* if the following are satisfied:

- for all $V \in \mathcal{V}$, $V \sqsubseteq W$;
- there exists $Q \sqsubseteq W$ such that $Q \perp V$ for all $V \in \mathcal{V}$.

Let $\mathfrak{D}$ denote the set of all container pairs. We will denote a pair $(W, \mathcal{V}) \in \mathfrak{D}$ by $D^{\mathcal{V}}_W$.

Let $\mathfrak{R} = \mathfrak{S} \cup \mathfrak{D}$. We will prove $\mathcal{X}$ has a hierarchically hyperbolic space structure with index set $\mathfrak{R}$. Since $(\mathcal{X}, \mathfrak{S})$ is an almost HHS, we can continue to use the spaces, projections, and relations for elements of $\mathfrak{S}$. Thus we only define new projections, relative projections, and relations when elements of $\mathfrak{D}$ are involved. If $D^{\mathcal{V}}_W \in \mathfrak{D}$, then the associated hyperbolic space, $\mathcal{C}D^{\mathcal{V}}_W$, will be a single point.

Projections: For $D^{\mathcal{V}}_W \in \mathfrak{D}$, the projection map is just the constant map to the single point in $\mathcal{C}D^{\mathcal{V}}_W$.

Nesting: Let $Q \in \mathfrak{S}$ and $D^{\mathcal{V}}_W, D^{\mathcal{R}}_T \in \mathfrak{D}$.

- Define $Q \sqsubseteq D^{\mathcal{V}}_W$ if $Q \sqsubseteq W$ in $\mathfrak{S}$ and $Q \perp V$ for all $V \in \mathcal{V}$.
- Define $D^{\mathcal{V}}_W \sqsubseteq Q$ if $W \sqsubseteq Q$ in $\mathfrak{S}$.
- Define $D^{\mathcal{V}}_W \sqsubseteq D^{\mathcal{R}}_T$ if $W \sqsubseteq T$ in $\mathfrak{S}$ and for all $R \in \mathcal{R}$ either $R \perp W$ or there exists $V \in \mathcal{V}$ with $R \sqsubseteq V$.

These definitions ensure $\sqsubseteq$ is still a partial order and maintain the $\sqsubseteq$–maximal element of $\mathfrak{S}$ as the $\sqsubseteq$–maximal element of $\mathfrak{R}$.

Since the hyperbolic spaces associated to elements of $\mathfrak{D}$ are points, define $\rho^Q_{D^{\mathcal{V}}_W} = \mathcal{C}D^{\mathcal{V}}_W$ for every $Q \in \mathfrak{R}$ and $D^{\mathcal{V}}_W \in \mathfrak{D}$ with $Q \sqsubsetneq D^{\mathcal{V}}_W$. The downwards relative projection $\rho^{D^{\mathcal{V}}_W}_Q : \mathcal{C}D^{\mathcal{V}}_W \to \mathcal{C}Q$ can be defined arbitrarily.

If $D^{\mathcal{V}}_W \in \mathfrak{D}$ and $Q \in \mathfrak{S}$ with $D^{\mathcal{V}}_W \sqsubseteq Q$, then $V \sqsubsetneq Q$ in $\mathfrak{S}$ for each $V \in \mathcal{V}$. Thus we define $\rho^{D^{\mathcal{V}}_W}_Q = \bigcup_{V \in \mathcal{V}} \rho^V_Q$. Lemma A.2 ensures that $\rho^{D^{\mathcal{V}}_W}_Q$ has diameter at most $4E$. In this case, we define $\rho^Q_{D^{\mathcal{V}}_W} : \mathcal{C}Q \to \mathcal{C}D^{\mathcal{V}}_W$ as the constant map to the single point in $\mathcal{C}D^{\mathcal{V}}_W$.

Finite complexity: First consider a nesting chain of the form $D^{\mathcal{V}_1}_W \sqsubsetneq D^{\mathcal{V}_2}_W \sqsubsetneq \ldots \sqsubsetneq D^{\mathcal{V}_n}_W$.

Claim A.5. *The length of* $D^{\mathcal{V}_1}_W \sqsubsetneq D^{\mathcal{V}_2}_W \sqsubsetneq \ldots \sqsubsetneq D^{\mathcal{V}_n}_W$ *is bounded above by* $E^2 + E$.

Proof. For each $V \in \bigcup_{i=1}^n \mathcal{V}_i$, we have $V \sqsubseteq W$ and hence $V \not\perp W$. As $D^{\mathcal{V}_{i-1}}_W \sqsubsetneq D^{\mathcal{V}_i}_W$ for each $i \in \{2, \ldots, n\}$, every element of $\mathcal{V}_i$ must therefore be nested into an element of $\mathcal{V}_{i-1}$. Denote the elements of $\mathcal{V}_i$ by $V^i_1, \ldots V^i_{k_i}$. Since each $\mathcal{V}_i$ is a pairwise orthogonal subset of

$\mathfrak{S}$, we have $k_i \leqslant E$ for each $i \in \{1,\dots,n\}$ by the bounded pairwise orthogonality axiom of an almost HHS (Definition 3.4). We define a $\mathcal{V}$*–nesting chain* to be a maximal chain of the form $V^m_{j_m} \sqsubseteq V^{m-1}_{j_{m-1}} \sqsubseteq \dots \sqsubseteq V^1_{j_1}$ for some $m \in \{1,\dots,n\}$ and $j_i \in \{1,\dots,k_i\}$, with $i \in \{1,\dots,m\}$. Since the elements of $\mathcal{V}_i$ are pairwise orthogonal for each $i \in \{1,\dots,n\}$, if $V^m_{j_m}$ is the $\sqsubseteq$–minimal element of a $\mathcal{V}$–nesting chain, then $V^m_{j_m}$ is nested into exactly one element of $\mathcal{V}_i$ for each $i \leqslant m$. This implies that each $\mathcal{V}$–nesting chain is determined by its $\sqsubseteq$–minimal element. Further, the set of $\sqsubseteq$–minimal elements of $\mathcal{V}$–nesting chains is pairwise orthogonal. By the bounded pairwise orthogonality axiom of an almost HHS, this implies there exist at most E $\mathcal{V}$–nesting chains.

In order for $D^{\mathcal{V}_i}_W \neq D^{\mathcal{V}_{i+1}}_W$, either $k_{i+1} < k_i$ or there exists $j_i \in \{1,\dots,k_i\}$, $j_{i+1} \in \{1,\dots,k_{i+1}\}$ such that $V^{i+1}_{j_{i+1}} \subsetneq V^i_{j_i}$. Thus, every step up the chain $D^{\mathcal{V}_1}_W \subsetneq D^{\mathcal{V}_2}_W \subsetneq \dots \subsetneq D^{\mathcal{V}_n}_W$ results in either a strict decrease in k_i (the cardinality of $\mathcal{V}_i$) to k_{i+1} (the cardinality of $\mathcal{V}_{i+1}$) or a strict step down one of the $\mathcal{V}$–nesting chains. Note that k_i may increase when we encounter a strict step down one of the $\mathcal{V}$–nesting chains, since multiple elements of $\mathcal{V}_{i+1}$ may be nested into the same element of $\mathcal{V}_i$. Such an increase in k_i corresponds to the nesting chain branching into multiple chains, which may only happen at most $E - k_1$ times, as there are at most E $\mathcal{V}$–nesting chains. Hence, the length of $D^{\mathcal{V}_1}_W \subsetneq D^{\mathcal{V}_2}_W \subsetneq \dots \subsetneq D^{\mathcal{V}_n}_W$ is bounded by $k_1 + (E - k_1) = E$ plus the total number of times a strict decrease can occur across all of the $\mathcal{V}$–nesting chains.

Each $\mathcal{V}$–nesting chain $V^m_{j_m} \sqsubseteq V^{m-1}_{j_{m-1}} \sqsubseteq \dots \sqsubseteq V^1_{j_1}$ contains at most E distinct elements of $\mathfrak{S}$ by the finite complexity of $\mathfrak{S}$. Bounded pairwise orthogonality implies there are at most E different $\mathcal{V}$–nesting chains, thus the number of steps of the chain $D^{\mathcal{V}_1}_W \subsetneq D^{\mathcal{V}_2}_W \subsetneq \dots \subsetneq D^{\mathcal{V}_n}_W$ where there is a strict decrease within one of the $\mathcal{V}$–nesting chains is at most E^2. This bounds the length of $D^{\mathcal{V}_1}_W \subsetneq D^{\mathcal{V}_2}_W \subsetneq \dots \subsetneq D^{\mathcal{V}_n}_W$ by $E^2 + E$. □

We now consider a nesting chain of the form $D^{\mathcal{V}_1}_{W_1} \subsetneq D^{\mathcal{V}_2}_{W_2} \subsetneq \dots \subsetneq D^{\mathcal{V}_n}_{W_n}$. In this case, $W_1 \sqsubseteq W_2 \sqsubseteq \dots \sqsubseteq W_n$, but not all of these nestings must be proper. Let $1 = i_1 < i_2 < \dots < i_k$ be the minimal subset of $\{1,\dots,n\}$ such that if $i_j \leqslant i < i_{j+1}$, then $W_{i_j} = W_i$. Thus $W_{i_1} \subsetneq W_{i_2} \subsetneq \dots \subsetneq W_{i_k}$, and $k \leqslant E$ by finite complexity of $\mathfrak{S}$. Claim A.5 established that $|i_j - i_{j+1}| \leqslant E^2 + E$, so $n \leqslant k(E^2+E) \leqslant E^3 + E^2$, that is, any $\subsetneq$–chain of elements of $\mathfrak{D}$ has length at most $E^3 + E^2$.

Finally, since any $\subsetneq$–chain of elements of $\mathfrak{R}$ can be partitioned into a $\subsetneq$–chain of elements of $\mathfrak{D}$ and a $\subsetneq$–chain of elements of $\mathfrak{S}$, any $\subsetneq$–chain in $\mathfrak{R}$ has length at most $E^3 + E^2 + E$.

Orthogonality: Two elements $D^{\mathcal{V}}_W, D^{\mathcal{R}}_T \in \mathfrak{D}$ are orthogonal if $W \perp T$ in $\mathfrak{S}$. Let $Q \in \mathfrak{S}$ and $D^{\mathcal{V}}_W \in \mathfrak{D}$. Define $Q \perp D^{\mathcal{V}}_W$ if, in $\mathfrak{S}$, either $W \perp Q$ or $Q \sqsubseteq V$ for some $V \in \mathcal{V}$. These definitions, plus the definition of nesting imply for all $W, V, Q \in \mathfrak{R}$, if $W \perp V$ and $Q \sqsubseteq V$, then $W \perp V$. We now verify that $\mathfrak{R}$ satisfies the container requirements of the orthogonality axiom.

Let $W, V \in \mathfrak{S}$ with $V \subsetneq W$ and $\{Q \in \mathfrak{R}_W : Q \perp V\} \neq \emptyset$, i.e., $(W, \{V\})$ is a container pair. In this case, the container of V in W for $\mathfrak{R}$ is $D^{\{V\}}_W$. We now show containers exist for situations involving elements of $\mathfrak{D}$. We split this into three subcases.

Case 1: $\mathbf{D^{\mathcal{V}}_W} \in \mathfrak{D}$ and $\mathbf{Q} \in \mathfrak{S}$ with $\mathbf{D^{\mathcal{V}}_W} \sqsubseteq \mathbf{Q}$. Since $(W, \mathcal{V})$ is a container pair, there exists $P \in \mathfrak{S}$ with $P \sqsubseteq W$ and $V \perp P$ for all $V \in \mathcal{V}$. Suppose that $D^{\mathcal{V}}_W$ requires a container in Q, that is, there is an element U of $\mathfrak{R}$ that is orthogonal to $D^{\mathcal{V}}_W$ and nested in Q. We verify that $(Q, \{P\})$ is a container pair and $D^{\{P\}}_Q$ is a container of $D^{\mathcal{V}}_W$ in Q.

If $T \in \mathfrak{S}$ with $T \perp D^{\mathcal{V}}_W$ and $T \sqsubseteq Q$, then $T \perp W$ or $T \sqsubseteq V$ for some $V \in \mathcal{V}$. In either case, we have $T \perp P$, so $(Q, \{P\})$ is a container pair and $T \sqsubseteq D^{\{P\}}_Q$. If $D^{\mathcal{R}}_T \in \mathfrak{D}$ with $D^{\mathcal{R}}_T \sqsubseteq Q$ and

$D_T^{\mathcal{R}} \perp D_W^{\mathcal{V}}$, then $T \perp W$ and $T \sqsubseteq Q$. Since $P \sqsubseteq W$, this implies $T \perp P$ and so $(Q, \{P\})$ is again a container pair, and $D_T^{\mathcal{R}} \sqsubseteq D_Q^{\{P\}}$.

Case 2: $\mathbf{D_W^{\mathcal{V}}, D_T^{\mathcal{R}} \in \mathfrak{D}}$ where $\mathbf{D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}}$. Since $(W, \mathcal{V})$ is a container pair, there exists $P \in \mathfrak{S}$ so that $P \sqsubseteq W$ and $P \perp V$ for all $V \in \mathcal{V}$. Since $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}$, it follows that for all $R \in \mathcal{R}$, either $R \perp W$ or there exists $V \in \mathcal{V}$ so that $R \sqsubseteq V$. In both cases, $R \perp P$. Thus $\mathcal{P} = \mathcal{R} \cup \{P\}$ is a pairwise orthogonal collection of elements of $\mathfrak{S}$. Suppose that $D_W^{\mathcal{V}}$ requires a container in $D_T^{\mathcal{R}}$, that is, there is an element U of $\mathfrak{R}$ that is orthogonal to $D_W^{\mathcal{V}}$ and nested in $D_T^{\mathcal{R}}$. We verify that $(T, \mathcal{P})$ is a container pair and $D_T^{\mathcal{P}} \subsetneq D_T^{\mathcal{R}}$ is a container for $D_W^{\mathcal{V}}$ in $D_T^{\mathcal{R}}$.

If $Q \in \mathfrak{S}$ satisfies $Q \sqsubseteq D_T^{\mathcal{R}}$ and $D_W^{\mathcal{V}} \perp Q$, then either $Q \perp W$ or $Q \sqsubseteq V$ for some $V \in \mathcal{V}$. In both cases, $Q \perp P$. Further, we must have $Q \sqsubseteq T$ and $Q \perp R$ for each $R \in \mathcal{R}$ as $Q \sqsubseteq D_T^{\mathcal{R}}$. Thus $(T, \mathcal{P})$ is a container pair and $Q \sqsubseteq D_T^{\mathcal{P}}$. On the other hand, if $D_Q^{\mathcal{Z}} \in \mathfrak{D}$ satisfies $D_Q^{\mathcal{Z}} \perp D_W^{\mathcal{V}}$ and $D_Q^{\mathcal{Z}} \sqsubseteq D_T^{\mathcal{R}}$, then $Q \perp W$, $Q \sqsubseteq T$, and for each $R \in \mathcal{R}$ either $R \perp Q$ or there exists $Z \in \mathcal{Z}$ with $R \sqsubseteq Z$. Since $(Q, \mathcal{Z})$ is a container pair, there exists $U \in \mathfrak{S}$ such that $U \sqsubseteq Q$ and $U \perp Z$ for all $Z \in \mathcal{Z}$. Since $Q \perp W$, we also have $U \perp P$ as $U \sqsubseteq Q$ and $P \sqsubseteq W$. For each $R \in \mathcal{R}$, either $R \perp Q$ or there exists $Z \in \mathcal{Z}$ with $R \sqsubseteq Z$. In both cases, $R \perp U$. Thus, U is orthogonal to all elements of $\mathcal{P} = \mathcal{R} \cup \{P\}$ and moreover $U \sqsubseteq Q \sqsubseteq T$, so $(T, \mathcal{P})$ is a container pair. Furthermore, $D_Q^{\mathcal{Z}} \sqsubseteq D_T^{\mathcal{P}} = D_T^{\mathcal{R} \cup \{P\}}$ since $D_Q^{\mathcal{Z}} \sqsubseteq D_T^{\mathcal{R}}$ and $P \perp Q$. We have therefore shown that $D_T^{\mathcal{P}}$ is a container for $D_W^{\mathcal{V}}$ in $D_T^{\mathcal{R}}$.

Case 3: $\mathbf{D_T^{\mathcal{R}} \in \mathfrak{D}}$ and $\mathbf{Q \in \mathfrak{S}}$ with $\mathbf{Q \sqsubseteq D_T^{\mathcal{R}}}$. This implies $\mathcal{Q} = \mathcal{R} \cup \{Q\}$ is a pairwise orthogonal set of elements of $\mathfrak{S}$. Further, suppose that Q requires a container in $D_T^{\mathcal{R}}$, that is, there is an element of $\mathfrak{R}$ that is orthogonal to Q and nested in $D_T^{\mathcal{R}}$. We verify that $(T, \mathcal{Q})$ is a container pair and $D_T^{\mathcal{Q}}$ is a container for Q in $D_T^{\mathcal{R}}$.

Suppose there exists $V \in \mathfrak{S}$ with $V \sqsubseteq D_T^{\mathcal{R}}$ and $V \perp Q$. Then $V \sqsubseteq T$ and V is orthogonal to all the elements of $\mathcal{R} \cup \{Q\}$. Thus $(T, \mathcal{Q})$ is a container pair, so $D_T^{\mathcal{Q}}$ exists and $V \sqsubseteq D_T^{\mathcal{Q}}$. Now suppose there exists $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}$ such that $D_W^{\mathcal{V}} \perp Q$. Since $(W, \mathcal{V})$ is a container pair, there exists $U \in \mathfrak{S}$ with $U \sqsubseteq W$ and U orthogonal to each element of $\mathcal{V}$. As $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}$, for each $R \in \mathcal{R}$ either $R \perp W$ or there exists $V \in \mathcal{V}$ such that $R \sqsubseteq V$. In both cases, $R \perp U$. Further, as $Q \perp D_W^{\mathcal{V}}$, we have $Q \perp W$ or $Q \sqsubseteq V$ for some $V \in \mathcal{V}$. In both cases, $Q \perp U$. Therefore U is orthogonal to every element of $\mathcal{Q}$, and moreover $U \sqsubseteq W \sqsubseteq T$ since $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}$. Thus $(T, \mathcal{Q})$ is a container pair and $U \sqsubseteq D_T^{\mathcal{Q}}$. Now, for each $R \in \mathcal{R}$, either $R \perp W$ or $R \sqsubseteq V$ for some $V \in \mathcal{V}$. Since $\mathcal{Q} = \mathcal{R} \cup \{Q\}$ and $Q \perp W$, this implies $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{Q}}$. Thus, $(T, \mathcal{Q})$ is a container pair and $D_T^{\mathcal{Q}}$ is a container for Q in $D_T^{\mathcal{R}}$.

Transversality: An element of $\mathfrak{R}$ is transverse to an element of $\mathfrak{D}$ whenever it is not nested or orthogonal. Since the hyperbolic spaces associated to elements of $\mathfrak{D}$ are points, we only need to define the relative projections from an element of $\mathfrak{D}$ to an element of $\mathfrak{S}$. Let $D_W^{\mathcal{V}} \in \mathfrak{D}$ and $Q \in \mathfrak{S}$ and suppose $D_W^{\mathcal{V}} \pitchfork Q$. This implies $W \not\perp Q$ and $W \not\sqsubseteq Q$. We define $\rho_Q^{D_W^{\mathcal{V}}}$ based on the $\mathfrak{S}$–relation between Q and the elements of $\mathcal{V}$.

- If $Q \perp V$ for all $V \in \mathcal{V}$, then $Q \not\sqsubseteq W$ as $Q \sqsubseteq W$ would imply $Q \sqsubseteq D_W^{\mathcal{V}}$. Thus we must have $Q \pitchfork W$, so we define $\rho_Q^{D_W^{\mathcal{V}}} = \rho_Q^W$.
- If $V \pitchfork Q$ or $V \subsetneq Q$ for some $V \in \mathcal{V}$, then ρ_Q^V exists and we define $\rho_Q^{D_W^{\mathcal{V}}}$ to be the union of all the ρ_Q^V for $V \in \mathcal{V}$ with $V \pitchfork Q$ or $V \subsetneq Q$. Lemma A.2 ensures $\rho_Q^{D_W^{\mathcal{V}}}$ has diameter at most $4E$ in this case.
- If $Q \sqsubseteq V$ for some V, then $Q \perp D_W^{\mathcal{V}}$ which contradicts $Q \pitchfork D_W^{\mathcal{V}}$, so this case does not occur.

Consistency: Since the only elements of $\mathfrak{R}$ whose associated spaces are not points are in $\mathfrak{S}$, the first two inequalities in the consistency axiom for $(\mathcal{X},\mathfrak{S})$ imply the same two inequalities for $(\mathcal{X},\mathfrak{R})$. To verify the final clause of the consistency axiom, we need to check that if $Q,R,T\in\mathfrak{R}$ such that $Q\subsetneq R$ with ρ^R_T and ρ^Q_T both defined, then $\mathsf{d}_T(\rho^Q_T,\rho^R_T)$ is uniformly bounded in terms of E. We can assume $T\in\mathfrak{S}$ as $\mathcal{C}T$ has diameter zero otherwise. We can further assume at least one of Q and R is an element of $\mathfrak{D}$, as we already have the consistency axiom for elements of $\mathfrak{S}$.

Case 1: $\mathbf{Q\subsetneq R\subsetneq T}$.

- Assume $Q\in\mathfrak{S}$ and $R=D^{\mathcal{V}}_W\in\mathfrak{D}$. Fix $V\in\mathcal{V}$. Since $D^{\mathcal{V}}_W=R\sqsubseteq T$ and $\rho^{D^{\mathcal{V}}_W}_T=\bigcup_{U\in\mathcal{V}}\rho^U_T$, we have $\rho^V_T\subseteq\rho^{D^{\mathcal{V}}_W}_T=\rho^R_T$. Since $V\perp Q$, Lemma A.2 says $\mathsf{d}_T(\rho^R_T,\rho^Q_T)\leqslant\mathsf{d}_T(\rho^V_T,\rho^Q_T)\leqslant 2E$.
- Assume $Q=D^{\mathcal{V}}_W\in\mathfrak{D}$ and $R\in\mathfrak{S}$. Fix $V\in\mathcal{V}$. In this case, $\rho^V_T\subseteq\rho^Q_T$ since $D^{\mathcal{V}}_W=Q\subsetneq T$. Since $D^{\mathcal{V}}_W=Q\subsetneq R$, we have $V\subsetneq W\subsetneq R$. Thus, the consistency axiom for $\mathfrak{S}$ says $\mathsf{d}_T(\rho^Q_T,\rho^R_T)\leqslant\mathsf{d}_T(\rho^V_T,\rho^R_T)\leqslant E$.
- Assume $Q=D^{\mathcal{V}}_W\in\mathfrak{D}$ and $R=D^{\mathcal{V}'}_{W'}\in\mathfrak{D}$. Thus $W\sqsubseteq W'\subsetneq T$ and consistency in $\mathfrak{S}$ implies $\mathsf{d}_T(\rho^W_T,\rho^{W'}_T)\leqslant E$. Fix $V\in\mathcal{V}$ and $V'\in\mathcal{V}'$. Consistency in $\mathfrak{S}$ also implies $\mathsf{d}_T(\rho^V_T,\rho^W_T)\leqslant E$ and $\mathsf{d}_T(\rho^{V'}_T,\rho^{W'}_T)\leqslant E$. Since $\rho^V_T\subseteq\rho^Q_T$ and $\rho^{V'}_T\subseteq\rho^R_T$, we have $\mathsf{d}_T(\rho^Q_T,\rho^R_T)\leqslant\mathsf{d}_T(\rho^V_T,\rho^{V'}_T)\leqslant\mathsf{d}_T(\rho^V_T,\rho^W_T)+\mathsf{diam}(\rho^W_T)+\mathsf{d}_T(\rho^W_T,\rho^{W'}_T)+\mathsf{diam}(\rho^{W'}_T)+\mathsf{d}_T(\rho^{W'}_T,\rho^{V'}_T)\leqslant 5E$.

Case 2: $\mathbf{Q\subsetneq R}$, $\mathbf{R\pitchfork T}$, and $\mathbf{Q\not\perp T}$. In this case we have either $Q\pitchfork T$ or $Q\subsetneq T$.

- Assume $Q\in\mathfrak{S}$ and $R=D^{\mathcal{V}}_W\in\mathfrak{D}$. Since $D^{\mathcal{V}}_W=R$ is transverse to T we cannot have $T\sqsubseteq V$ for any $V\in\mathcal{V}$ (this would imply $D^{\mathcal{V}}_W\perp T$). If $V\perp T$ for all $V\in\mathcal{V}$, then $W\pitchfork T$ (as shown in the proof of transversality) and $\rho^R_T=\rho^{D^{\mathcal{V}}_W}_T=\rho^W_T$. Since $Q\sqsubseteq R=D^{\mathcal{V}}_W$, we have $Q\sqsubseteq W$ and consistency in $\mathfrak{S}$ implies $\mathsf{d}_T(\rho^Q_T,\rho^R_T)=\mathsf{d}_T(\rho^Q_T,\rho^W_T)\leqslant E$. If instead there exists $V\in\mathcal{V}$ so that $V\pitchfork T$ or $V\subsetneq T$, then $\rho^V_T\subseteq\rho^{D^{\mathcal{V}}_W}_T=\rho^R_T$. Since $Q\sqsubseteq R=D^{\mathcal{V}}_W$, $Q\perp V$ and Lemma A.2 gives $\mathsf{d}_T(\rho^Q_T,\rho^R_T)\leqslant\mathsf{d}_T(\rho^Q_T,\rho^V_T)\leqslant 2E$.
- Assume $Q=D^{\mathcal{V}}_W\in\mathfrak{D}$ and $R\in\mathfrak{S}$. As before, $T\not\sqsubseteq V$ for all $V\in\mathcal{V}$. First assume there exists $V\in\mathcal{V}$ so that $V\pitchfork T$ or $V\subsetneq T$. This occurs when either $D^{\mathcal{V}}_W=Q\subsetneq T$ or $Q\pitchfork T$ and not every element of $\mathcal{V}$ is orthogonal to T. In both cases, $\rho^V_T\subseteq\rho^{D^{\mathcal{V}}_W}_T=\rho^Q_T$ and consistency in $\mathfrak{S}$ implies $\mathsf{d}_T(\rho^Q_T,\rho^R_T)\leqslant\mathsf{d}_T(\rho^V_T,\rho^R_T)\leqslant 2E$ because $V\sqsubseteq W\subsetneq R$. Now assume $T\perp V$ for all $V\in\mathcal{V}$. This can only occur when $D^{\mathcal{V}}_W=Q$ is transverse to T. In this case, $W\pitchfork T$ and $\rho^Q_T=\rho^{D^{\mathcal{V}}_W}_T=\rho^W_T$. Since $W\sqsubseteq R$, consistency in $\mathfrak{S}$ implies $\mathsf{d}_T(\rho^R_T,\rho^Q_T)=\mathsf{d}_T(\rho^R_T,\rho^W_T)\leqslant E$.
- Assume $Q=D^{\mathcal{V}}_W\in\mathfrak{D}$ and $R=D^{\mathcal{V}'}_{W'}\in\mathfrak{D}$. As before, $T\not\sqsubseteq V$ for all $V\in\mathcal{V}\cup\mathcal{V}'$. If $\rho^R_T=\rho^{W'}_T$, then we have the first case of transversality, that is, $W'\pitchfork T$ and $V'\perp T$ for all $V'\in\mathcal{V}$. Thus, if $\rho^R_T=\rho^{W'}_T$, then the result reduces to the previous bullet, replacing R with W'. We can therefore assume $\rho^R_T\neq\rho^{W'}_T$, meaning we have the second case of transversality where there exists $V'\in\mathcal{V}'$ so that V' is either transverse to or properly nested into T.

  Suppose $\rho^Q_T\neq\rho^W_T$ too. This implies there also exists $V\in\mathcal{V}$ so that V is either transverse to or properly nested into T. Furthermore, $\rho^V_T\subseteq\rho^Q_T$ and $\rho^{V'}_T\subseteq\rho^R_T$. Now, $D^{\mathcal{V}}_W\sqsubseteq D^{\mathcal{V}'}_{W'}$ implies $V'\perp W$ or V' is nested into an element of $\mathcal{V}$. If $V'\perp W$, then $V\perp V'$ and Lemma A.2 implies $\mathsf{d}_T(\rho^Q_T,\rho^R_T)\leqslant\mathsf{d}_T(\rho^V_T,\rho^{V'}_T)\leqslant 2E$. If V' is nested into an element of $\mathcal{V}$, then either $V'\sqsubseteq V$ or $V'\perp V$ since $\mathcal{V}$ is a pairwise orthogonal subset of $\mathfrak{S}$. By applying consistency in $\mathfrak{S}$ when $V' \sqsubseteq V$ or Lemma A.2 when $V' \perp V$, we have $\mathsf{d}_T(\rho_T^Q, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^V, \rho_T^{V'}) \leqslant 2E$.

Now suppose $\rho_T^Q = \rho_T^W$. Then $D_W^{\mathcal{V}} \sqsubseteq D_{W'}^{\mathcal{V}'}$ implies $V' \perp W$ or V' is nested into W. Applying Lemma A.2 if $V' \perp W$, or consistency in $\mathfrak{S}$ if $V' \sqsubseteq W$, we again obtain $\mathsf{d}_T(\rho_T^Q, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^W, \rho_T^{V'}) \leqslant 2E$.

Uniqueness, bounded geodesic image, large links: Since the only elements of $\mathfrak{R}$ whose associated spaces are not points are in $\mathfrak{S}$, these axioms for $(\mathcal{X}, \mathfrak{R})$ follow from the fact that they hold in $(\mathcal{X}, \mathfrak{S})$.

Partial realization: Let $T_1, \ldots, T_n$ be pairwise orthogonal elements of $\mathfrak{R}$, and let $p_i \in \mathcal{C}T_i$ for each $i \in \{1, \ldots, n\}$. Without loss of generality, assume $T_1, \ldots, T_k \in \mathfrak{S}$ and $T_{k+1}, \ldots, T_n \in \mathfrak{D}$ where $k \in \{0, \ldots, n\}$. If $k = 0$ (resp. $k = n$), then each $T_i \in \mathfrak{D}$ (resp. $T_i \in \mathfrak{S}$).

For $i \in \{k+1, \ldots, n\}$, let $T_i = D_{W_i}^{\mathcal{V}_i}$ and let q_i be any point in $\rho_{W_i}^{D_{W_i}^{\mathcal{V}_i}} \subseteq \mathcal{C}W_i$. Since $T_1, \ldots, T_n$ are pairwise orthogonal, it follows that $W_{k+1}, \ldots, W_n$ are pairwise orthogonal too, and for each $j \in \{1, \ldots, k\}$, T_j is either nested into an element of some $\mathcal{V}_{i_j}$ or orthogonal to all $W_{k+1}, \ldots, W_n$. Without loss of generality, assume that $T_1, \ldots, T_l$ are nested into elements of $\mathcal{V}_{m+1} \cup \cdots \cup \mathcal{V}_n$ and $T_{l+1}, \ldots, T_k, W_{k+1}, \ldots, W_n$ are pairwise orthogonal, where $l \leqslant k$, $m \leqslant n$, and $n - m \leqslant l$. If $l = 0$, then $n = m$ and each T_j is orthogonal to every W_i. Otherwise, for each $j \in \{1, \ldots, l\}$, T_j is nested in some W_i for $i \in \{m+1, \ldots, n\}$. In both cases, $T_1, \ldots, T_k, W_{k+1}, \ldots, W_m$ are pairwise orthogonal elements of $\mathfrak{S}$. We can therefore use the partial realization axiom in $\mathfrak{S}$ on the points $p_1, \ldots, p_k, q_{k+1}, \ldots, q_m$ to produce a point $x \in \mathcal{X}$ with the following properties:

(1) $\mathsf{d}_{T_i}(x, p_i) \leqslant E$ for $i \in \{1, \ldots, k\}$;
(2) $\mathsf{d}_{W_i}(x, q_i) \leqslant E$ for $i \in \{k+1, \ldots, m\}$;
(3) for all $i \in \{1, \ldots, k\}$ if $Q \pitchfork T_i$ or $T_i \sqsubsetneq Q$, then $\mathsf{d}_Q(x, \rho_Q^{T_i}) \leqslant E$;
(4) for all $i \in \{k+1, \ldots, m\}$ if $Q \pitchfork W_i$ or $W_i \sqsubsetneq Q$, then $\mathsf{d}_Q(x, \rho_Q^{W_i}) \leqslant E$.

Now, for $Q \in \mathfrak{S}$, define $b_Q \in \mathcal{C}Q$ as follows. Let $\mathcal{V} = \bigcup_{i=k+1}^{n} \mathcal{V}_i$ and $\mathcal{V}_Q = \{V \in \mathcal{V} : V \pitchfork Q \text{ or } V \sqsubsetneq Q\}$. If $\mathcal{V}_Q \neq \emptyset$, then define b_Q to be any point in $\bigcup_{V \in \mathcal{V}_Q} \rho_Q^V$. Since $\mathcal{V}$ is a collection of pairwise orthogonal elements of $\mathfrak{S}$, the diameter of $\bigcup_{V \in \mathcal{V}_Q} \rho_Q^V$ is at most $2E$ by Lemma A.2. If either $Q \sqsubseteq V$ for some $V \in \mathcal{V}$ or $Q \perp V$ for all $V \in \mathcal{V}$ then define $b_Q = \pi_Q(x)$. Since $\mathcal{V}$ is a collection of pairwise orthogonal elements of $\mathfrak{S}$, these two cases encompass all elements of $\mathfrak{S}$.

Claim A.6. *The tuple $(b_Q)_{Q \in \mathfrak{S}}$ is $3E$–consistent.*

Proof. Let $R, Z \in \mathfrak{S}$. Recall that if $b_Z = \pi_Z(x)$ and $b_R = \pi_R(x)$, then the E–consistency inequalities for b_R and b_Z are satisfied by the consitency axiom of $(\mathcal{X}, \mathfrak{S})$. Thus we can assume that there exists $V \in \mathcal{V}$ so that either $V \sqsubsetneq Z$ or $V \pitchfork Z$. Fix $V \in \mathcal{V}$ so that $b_Z \in \rho_Z^V$. We need to verify the consistency inequalities when $R \pitchfork Z$, $R \sqsubsetneq Z$, and $Z \sqsubsetneq R$.

Consistency when $\mathbf{R} \pitchfork \mathbf{Z}$: Assume $R \pitchfork Z$. If $R \perp V$, $V \sqsubseteq R$, or $R \sqsubseteq V$ then either Lemma A.2 or consistency in $\mathfrak{S}$ implies $\mathsf{d}_Z(\rho_Z^V, \rho_Z^R) \leqslant 2E$. Since $b_Z \in \rho_Z^V$, we have $\mathsf{d}_Z(b_Z, \rho_Z^R) \leqslant 3E$. Now suppose $R \pitchfork V$ so that $\mathcal{V}_R$ is non-empty. In this case, $b_R \in \bigcup_{U \in \mathcal{V}_R} \rho_R^U$ and so b_R is within $2E$ of ρ_R^V. Now, if $\mathsf{d}_Z(b_Z, \rho_Z^R) > 3E$, then $\mathsf{d}_Z(\rho_Z^V, \rho_Z^R) > 2E$. Thus ρ–consistency (Lemma A.4) implies $\mathsf{d}_R(\rho_R^V, \rho_R^Z) \leqslant E$. It follows that $\mathsf{d}_R(b_R, \rho_R^Z) \leqslant 3E$ by the triangle inequality.

Consistency when $\mathbf{R} \sqsubsetneq \mathbf{Z}$: Assume $R \sqsubsetneq Z$. As before, if $R \perp V$, $V \sqsubseteq R$, or $R \sqsubseteq V$ then $\mathsf{d}_Z(\rho_Z^V, \rho_Z^R) \leqslant 2E$ and we have $\mathsf{d}_Z(b_Z, \rho_Z^R) \leqslant 3E$. Thus, we can assume $R \pitchfork V$ so that b_R is within $2E$ of ρ_R^V. Now, if $\mathsf{d}_Z(b_Z, \rho_Z^R) > 3E$, then $\mathsf{d}_Z(\rho_Z^V, \rho_Z^R) > 2E$, and ρ–consistency implies $\mathsf{diam}(\rho_R^V \cup \rho_R^Z(\rho_Z^V)) \leqslant E$. However, this implies $\mathsf{diam}(b_R \cup \rho_R^Z(b_Z)) \leqslant 3E$ since $b_Z \in \rho_Z^V$ and $\mathsf{d}_R(b_R, \rho_R^V) \leqslant 2E$.

Consistency when $\mathbf{Z} \subsetneq \mathbf{R}$: Assume $Z \subsetneq R$. If R is orthogonal to all elements of $\mathcal{V}$, then $R \perp V$ implies $V \perp Z$ which contradicts the assumption that $V \subsetneq Z$ or $V \pitchfork Z$. On the other hand, if there exists $V' \in \mathcal{V}$ so that $R \sqsubseteq V'$, then either $R \perp V$ (if $V' \perp V$) or $R \sqsubseteq V$ (if $V' = V$). But this implies either $V \perp Z$ or $Z \subsetneq V$, both of which give a contradiction if $V \pitchfork Z$ or $V \subsetneq Z$. There must therefore be an element of $\mathcal{V}$ that is either properly nested in or transverse to R, and we can repeat the same argument as in the previous case, switching the roles of R and Z. □

Let $y \in \mathcal{X}$ be the point produced by applying the realization theorem (Theorem A.3) in $\mathfrak{S}$ to the tuple (b_Q). We claim y is a partial realization point for $p_1, \ldots, p_n$ in $\mathfrak{R}$. Since $\mathcal{C}D^{\mathcal{V}_i}_{W_i}$ is a single point, y satisfies the first requirement of the partial realization axiom in $\mathfrak{R}$ for $p_{k+1}, \ldots, p_n$. For $i \leqslant k$, T_i is either nested into an element of $\mathcal{V}_{m+1} \cup \cdots \cup \mathcal{V}_n$ or orthogonal to all $W_{k+1}, \ldots, W_n$. This implies T_i is either nested into an element of $\mathcal{V}$ or orthogonal to all elements of $\mathcal{V}$. In both cases $b_{T_i} = \pi_{T_i}(x)$, and we have that $\pi_{T_i}(y)$ is uniformly close to $\pi_{T_i}(x)$, which is in turn E–close to p_i.

Now, let $Q \in \mathfrak{S}$ with $Q \pitchfork T_i$ or $T_i \subsetneq Q$ for some $i \in \{1, \ldots, n\}$. We verify $\mathsf{d}_Q(y, \rho^{T_i}_Q)$ is uniformly bounded when $i \leqslant k$ and $i > k$ separately.

Assume $i \leqslant k$, so that $T_i \in \mathfrak{S}$. If $i \leqslant k$ and $b_Q = \pi_Q(x)$, then $\mathsf{d}_Q(y, \rho^{T_i}_Q)$ is bounded by item (3). If $i \leqslant k$ and $b_Q \neq \pi_Q(x)$, then $b_Q \in \rho^V_Q$ for some $V \in \mathcal{V}$ and T_i is either orthogonal to or nested into V. If $T_i \perp V$ then $\mathsf{d}_Q(b_Q, \rho^{T_i}_Q) \leqslant 3E$ by Lemma A.2. If $T_i \subsetneq V$ then $\mathsf{d}_Q(b_Q, \rho^{T_i}_Q) \leqslant 2E$ by consistency. The result then follows from the triangle inequality since $\pi_Q(y)$ is uniformly close to b_Q.

Now assume $i > k$, so that $T_i = D^{\mathcal{V}_i}_{W_i} \in \mathfrak{D}$. If $D^{\mathcal{V}_i}_{W_i} \subsetneq Q$, then $\rho^V_Q \subseteq \rho^{D^{\mathcal{V}_i}_{W_i}}_Q$ for all $V \in \mathcal{V}_i$. Since b_Q is within $2E$ of any ρ^V_Q for $V \in \mathcal{V}_i$, this bounds $\mathsf{d}_Q(y, \rho^{D^{\mathcal{V}_i}_{W_i}}_Q)$ uniformly. On the other hand, if $D^{\mathcal{V}_i}_{W_i} \pitchfork Q$, then either $Q \perp V$ for all $V \in \mathcal{V}_i$ or there exists $V \in \mathcal{V}_i$ so that $V \pitchfork Q$ or $V \subsetneq Q$. In the latter case, $\rho^V_Q \subseteq \rho^{D^{\mathcal{V}_i}_{W_i}}_Q$ and we are finished since b_Q is within $2E$ of ρ^V_Q, giving a uniform bound on the distance from $\pi_Q(y)$ to $\rho^{D^{\mathcal{V}_i}_{W_i}}_Q$. In the former case, we must have $W_i \pitchfork Q$ and $\rho^{D^{\mathcal{V}_i}_{W_i}}_Q$ is equal to $\rho^{W_i}_Q$. If $b_Q = \pi_Q(x)$ than we are done by item (4). Otherwise, there exists $V' \in \mathcal{V} - \mathcal{V}_i$ so that $V' \pitchfork Q$ or $V' \subsetneq Q$ and $b_Q \in \rho^{V'}_Q$. Since $V' \perp W_i$, it follows that $\rho^{V'}_Q$ is within $2E$ of $\rho^{W_i}_Q$. Thus b_Q, and hence $\pi_Q(y)$, is uniformly close to $\rho^{W_i}_Q = \rho^{D^{\mathcal{V}_i}_{W_i}}_Q$. This concludes the proof of Theorem A.1. □

Remark A.7. We say G is an almost HHG if there exists an almost HHS $(\mathcal{X}, \mathfrak{S})$ such that G and $(\mathcal{X}, \mathfrak{S})$ satisfy the definition of a hierarchically hyperbolic group where 'HHS' is replaced with 'almost HHS'. The above proof shows that if $(G, \mathfrak{S})$ is an almost HHG, then the structure $\mathfrak{R}$ from Theorem A.1 is an HHG structure for G.

The following corollary gives criteria for the HHS structure from Theorem A.1 to have unbounded products. This is the version of Theorem A.1 that is applied in Theorem 3.7 to prove that every hierarchically hyperbolic space with the bounded domain dichotomy admits an HHS structure with unbounded products.

Corollary A.8. *Let $(\mathcal{X}, \mathfrak{T})$ be an almost HHS with the bounded domain dichotomy. If for every non-$\sqsubseteq$–maximal domain $V \in \mathfrak{T}$, there exist $W, Q \in \mathfrak{T}$ so that $W \sqsubseteq V$, $Q \perp V$, and* $\mathsf{diam}(\mathcal{C}W) = \mathsf{diam}(\mathcal{C}Q) = \infty$, *then the HHS structure $\mathfrak{R}$ obtained by applying Theorem A.1 to $\mathfrak{T}$ has unbounded products.*

Proof. Assume for every non-$\sqsubseteq$–maximal domain $V \in \mathfrak{T}$, there exist $W, Q \in \mathfrak{T}$ so that $W \sqsubseteq V$, $Q \perp V$ and $\mathsf{diam}(\mathcal{C}W) = \mathsf{diam}(\mathcal{C}Q) = \infty$. Let $\mathfrak{R}$ be the HHS structure obtained from $\mathfrak{T}$ using Theorem A.1. If $V \in \mathfrak{T}$ and V is not $\sqsubseteq$–maximal, then the above property implies that $\mathbf{F}_V$ and $\mathbf{E}_V$ are both infinite diameter. Thus, we need only verify unbounded products for elements of $\mathfrak{R} - \mathfrak{T}$. Using the notation of Theorem A.1, let $D = D_W^{\mathcal{V}} \in \mathfrak{R} - \mathfrak{T}$ and assume $\mathsf{diam}(\mathbf{F}_D) = \infty$. Now, $V \perp D_W^{\mathcal{V}}$ for all $V \in \mathcal{V}$, and by construction of $\mathfrak{T}$, there exists $Q \in \mathfrak{T}$ so that $Q \sqsubseteq V$ and $\mathsf{diam}(\mathcal{C}Q) = \infty$. Since $Q \perp D_W^{\mathcal{V}}$, this implies $\mathsf{diam}(\mathbf{E}_D) = \infty$. Therefore $(\mathcal{X}, \mathfrak{R})$ is an HHS with unbounded products. □

References

[Abb16] Carolyn R. Abbott. Not all finitely generated groups have universal acylindrical actions. *Proc. Amer. Math. Soc.*, 144(10):4151–4155, 2016.

[ABO19] Carolyn Abbott, Sahana H. Balasubramanya, and Denis Osin. Hyperbolic structures on groups. *Algebr. Geom. Topol.*, 19(4):1747–1835, 2019.

[ACGH17] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. Characterizations of Morse quasi-geodesics via superlinear divergence and sublinear contraction. *Doc. Math.*, 22:1193–1224, 2017.

[ADT17] Tarik Aougab, Matthew Gentry Durham, and Samuel J. Taylor. Pulling back stability with applications to $\mathrm{Out}(F_n)$ and relatively hyperbolic groups. *J. Lond. Math. Soc. (2)*, 96(3):565–583, 2017.

[BBKL16] Bestvina, Bromberg, Kent, and Leininger. Undistorted purely pseudo-Anosov groups. *arXiv:1608.01583*, 2016.

[Beh06] J. Behrstock. Asymptotic geometry of the mapping class group and Teichmüller space. *Geometry & Topology*, 10:1523–1578, 2006.

[BHS17a] J. Behrstock, M.F. Hagen, and A. Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. London Math. Soc.*, 114(5):890–926, 2017.

[BHS17b] Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces I: curve complexes for cubical groups. *Geometry & Topology*, 21:1731–1804, 2017.

[BHS17c] Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *arXiv:1704.04271*, 2017.

[BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299(2):257–338, 2019.

[Bow08] Brian H. Bowditch. Tight geodesics in the curve complex. *Inventiones mathematicae*, 171(2):281–300, 2008.

[BR] F. Berlai and B. Robbio. A refined combination theorem for hierarchically hyperbolic groups. *Groups Geometry and Dynamics.* To appear.

[Bra99] N. Brady. Branched coverings of cubical complexes and subgroups of hyperbolic groups. *J. London Math. Soc. (2)*, 60(2):461–480, 1999.

[CD19] Matthew Cordes and Matthew Gentry Durham. Boundary convex cocompactness and stability of subgroups of finitely generated groups. *Int. Math. Res. Not. IMRN*, (6):1699–1724, 2019.

[CH17] Matthew Cordes and David Hume. Stability and the Morse boundary. *J. Lond. Math. Soc. (2)*, 95(3):963–988, 2017.

[CM19] Christopher H. Cashen and John M. Mackay. A metrizable topology on the contracting boundary of a group. *Trans. Amer. Math. Soc.*, 372(3):1555–1600, 2019.

[Cor17] Matthew Cordes. Morse boundaries of proper geodesic metric spaces. *Groups Geom. Dyn.*, 11(4):1281–1306, 2017.

[CS15] Ruth Charney and Harold Sultan. Contracting boundaries of CAT(0) spaces. *J. Topol.*, 8(1):93–117, 2015.

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

[DHS17] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.

[DMS10] Cornelia Druţu, Shahar Mozes, and Mark Sapir. Divergence in lattices in semisimple Lie groups and graphs of groups. *Trans. Amer. Math. Soc.*, 362(5):2451–2505, 2010.

[DT15] Matthew Gentry Durham and Samuel J. Taylor. Convex cocompactness and stability in mapping class groups. *Algebr. Geom. Topol.*, 15(5):2839–2859, 2015.

[FM02] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152 (electronic), 2002.

[Hag14] Mark F. Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.*, 7(2):385–418, 2014.

[Ham] U. Hamenstädt. Word hyperbolic extensions of surface groups. Preprint, arXiv:0505244.

[HS16] M.F. Hagen and T. Susse. On hierarchical hyperbolicity of cubical groups. arXiv:1609.01313, 2016.

[KK14] Sang-Hyun Kim and Thomas Koberda. The geometry of the curve graph of a right-angled Artin group. *International Journal of Algebra and Computation*, 24(02):121–169, 2014.

[KL08] Kent and Leininger. Shadows of mapping class groups: capturing convex cocompactness. *Geom. Funct. Anal.*, 18(4):1270–1325, 2008.

[KMT17] Thomas Koberda, Johanna Mangahas, and Samuel J. Taylor. The geometry of purely loxodromic subgroups of right-angled Artin groups. *Trans. Amer. Math. Soc.*, 369(11):8179–8208, 2017.

[Min96] Yair N. Minsky. Quasi-projections in Teichmüller space. *J. Reine Angew. Math.*, 473:121–136, 1996.

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

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

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

[Rus19] Jacob Russell. From hierarchical to relative hyperbolicity. *arXiv preprint arXiv:1905.12489*, 2019.

[Sel97] Z. Sela. Acylindrical accessibility for groups. *Invent. Math.*, 129(3):527–565, 1997.

[Sis16] Alessandro Sisto. Quasiconvexity of hyperbolically embedded subgroups. *Mathematische Zeitschrift*, 283:649–658, 2016.

[Sul14] Harold Sultan. Hyperbolic quasi-geodesics in CAT(0) spaces. *Geom. Dedicata*, 169:209–224, 2014.

[Tra] Hung Cong Tran. Purely loxodromic subgroups in right-angled Coxeter groups. arXiv:1703.09032.

[TT16] Samuel J. Taylor and Giulio Tiozzo. Random extensions of free groups and surface groups are hyperbolic. *Int. Math. Res. Not. IMRN*, (1):294–310, 2016.

Columbia University, New York, New York, USA
Email address: abbott@math.columbia.edu

Lehman College and The Graduate Center, CUNY, New York, New York, USA
Email address: jason.behrstock@lehman.cuny.edu

University of California, Riverside, Riverside, California, USA
Email address: matthew.durham@ucr.edu