QUASIFLATS IN HIERARCHICALLY HYPERBOLIC SPACES

JASON BEHRSTOCK, MARK F HAGEN, AND ALESSANDRO SISTO

Abstract. The rank of a hierarchically hyperbolic space is the maximal number of unbounded factors of standard product regions; this coincides with the maximal dimension of a quasiflat for hierarchically hyperbolic groups. Several noteworthy examples for which the rank coincides with familiar quantities include: the dimension of maximal Dehn twist flats for mapping class groups, the maximal rank of a free abelian subgroup for right-angled Coxeter groups and right-angled Artin groups (in the latter this can also be observed as the clique number of the defining graph), and, for the Weil–Petersson metric the rank is the integer part of half the complex dimension of Teichmüller space.

We prove that, in a hierarchically hyperbolic space, any quasiflat of dimension equal to the rank lies within finite distance of a union of standard orthants (under a very mild condition on the HHS satisfied by all natural examples). This resolves outstanding conjectures when applied to a number of different groups and spaces. In the case of the mapping class group we verify a conjecture of Farb, for Teichmüller space we answer a question of Brock, and in the context of CAT(0) cubical groups our result proves a folk conjecture, novel special cases of the cubical case include right-angled Coxeter groups, and others.

An important ingredient in the proof, which we expect will have other applications, is our proof that the *hull* of any finite set in an HHS is quasi-isometric to a cube complex of dimension equal to the rank (if the HHS is a CAT(0) cube complex, the rank can be lower than the dimension of the space).

We deduce a number of applications of these results; for instance we show that any quasi-isometry between HHS induces a quasi-isometry between certain *factored spaces*, which are simpler HHS. This allows one, for example, to distinguish quasi-isometry classes of right-angled Artin/Coxeter groups.

Another application of our results is to quasi-isometric rigidity. Our tools in many cases allow one to reduce the problem of quasi-isometric rigidity for a given hierarchically hyperbolic group to a combinatorial problem. As a template, we give a new proof of quasi-isometric rigidity of mapping class groups, using simpler combinatorial arguments than in previous proofs.

Contents

Introduction 2
1. Background 7
2. Cubulation of hulls 16
3. Quasiflats and asymptotic cones 23
4. Orthants and quasiflats 25
5. Induced maps on hinges: mapping class group rigidity 30
6. Factored spaces 34
References 36

Date: February 15, 2018.

Introduction

A classical result of Morse shows that in a hyperbolic space quasigeodesics lie close to geodesics [Mor24]. This raises the question of what constraints exist on the geometry of quasiflats in more general non-positively curved spaces. A key step in proving Mostow Rigidity is proving that an equivariant quasi-isometry of a symmetric space sends each flat to within a bounded neighborhood of a flat [Mos73]. Unlike the case of quasigeodesics in hyperbolic space, in general, a quasiflat need not lie close to any one flat. Generalizing Mostow's result, in a higher-rank symmetric space an arbitrary quasiflat must lie close to a finite number of flats [EF97, KL97b]. This result can be used to prove quasi-isometric rigidity for uniform lattices in higher-rank symmetric spaces [KL97b], see also [EF97].

In this paper, we control the structure of quasiflats in a broad class of spaces and groups with a property called *hierarchical hyperbolicity* [BHS17b, BHS15, BHS17a]. This class effectively captures the negative curvature phenomena visible in many important groups and spaces, including mapping class groups, right-angled Artin groups, CAT(0) cube complexes with geometric group actions, most 3–manifold groups, Teichmüller space (in any of the standard metrics), etc.

The class of hierarchically hyperbolic spaces also includes many examples not on the preceding list, and in fact the definition gives great flexibility in building new hierarchically hyperbolic spaces from old. In particular, trees of hierarchically hyperbolic spaces satisfying natural constraints (and thus many graphs of hierarchically hyperbolic groups) are again hierarchically hyperbolic [BHS15]. Groups that are hyperbolic relative to hierarchically hyperbolic groups are again hierarchically hyperbolic [BHS15]. It is shown in [BHS17a] that suitable small-cancellation quotients of hierarchically hyperbolic groups are again hierarchically hyperbolic.

This article establishes a deep relationship between these examples: in particular, we establish the surprising fact that these spaces all admit a very strong local approximation by CAT(0) cube complexes (Theorem F). This bridge allows us to use cubical techniques in new settings such as the mapping class group. Even for CAT(0) cube complexes our approximation provides new information, as the approximating space is typically lower dimensional, which is essential to the argument. In this sense, these techniques are genuinely intrinsic to the category of hierarchically hyperbolic space.

Formal definitions and relevant properties of hierarchically hyperbolic spaces (HHS) will be given below in Section 1. For now, we recall that a *hierarchically hyperbolic space* consists of: a space, $\mathcal{X}$; an *index set*, $\mathfrak{S}$, for which each $U \in \mathfrak{S}$ is associated with a hyperbolic space $\mathcal{C}U$; and, some maps and relations between elements of the index set.

Before stating the main theorem, we informally recall a few facts about the geometry of HHS. Any HHS $\mathcal{X}$ contains certain *standard product regions*, in which each of the (boundedly many) factors is an HHS itself. In mapping class groups, these are products of mapping class groups of pairwise disjoint subsurfaces, and in cube complexes these are certain convex subcomplexes that split as products. Pairs of points in $\mathcal{X}$ can be joined by particularly well-behaved quasigeodesics called *hierarchy paths*, and similarly we have well-behaved quasigeodesic rays called *hierarchy rays*. Given a standard product region P, and a hierarchy ray in each of the k factors of P, the product of the k hierarchy rays $[0, \infty) \to \mathcal{X}$ is a quasi-isometric embedding $[0, \infty)^k \to \mathcal{X}$ which we call a *standard orthant*.

The *rank* ν of an HHS is the largest possible number of factors in a standard product region, each of whose factors is unbounded. (Equivalently, it is the maximal integer so that there exist pairwise *orthogonal* $U_1, \ldots, U_\nu \in \mathfrak{S}$ for which each $\mathcal{C}U_i$ is unbounded.) We will impose a mild technical assumption on our spaces, which we call being *asymphoric*; this condition is satisfied by the motivating examples of HHS, including all hierarchically hyperbolic groups. Under this condition, Theorem 1.14 implies that the rank is a quasi-isometry invariant.

Theorem A (Quasiflats Theorem for HHS)**.** *Let $\mathcal{X}$ be an HHG of rank ν. Let $f\colon \mathbb{R}^\nu \to \mathcal{X}$ be a quasi-isometric embedding. Then there exist standard orthants $Q_i \subseteq \mathcal{X}$, $i = 1, \dots, k$, so that $\mathsf{d}_{haus}(f(\mathbb{R}^\nu), \cup_{i=1}^k Q_i) < \infty$. More generally, the same result holds for any space $\mathcal{X}$ which is an asymphoric HHS of rank ν.*

We now give a few immediate applications of this theorem.

Mapping class groups are hierarchically hyperbolic, by [BHS15, Theorem 11.1]. Theorem A applied to this case resolves a conjecture of Farb. Outside of the hyperbolic cases, this question was completely open.

Corollary B (Farb Conjecture: Quasiflats theorem for mapping class groups)**.** *Any top-dimensional quasiflat in the mapping class group is uniformly close to a finite union of standard flats.*

Outside of groups, an interesting application of Theorem A is to the Weil-Petersson metric on Teichmüller space, which was proven to be an asymphoric HHS in [BHS17b, Theorem G]. Brock asked whether every top-dimensional quasiflat in the Weil-Petersson metric is a bounded distance from a finite number of top-dimensional flats [Bro02, Question 5.3]. From Theorem A we obtain the following, which completely resolves his question in the affirmative.

Corollary C (Affirmation of Brock's Questions: Quasiflats theorem for Weil-Petersson metric)**.** *Any top-dimensional quasiflat in the Weil-Petersson metric on Teichmüller space is uniformly close to a finite union of standard flats.*

The only previously known cases of Brock's question were: in the rank one cases, where the space is hyperbolic [BF06]; and, in the three rank two cases, where the space is relatively hyperbolic, [BM08, Theorem 3]. There also existed partial results about flats being locally contained in linear size neighborhoods of standard flats, e.g., [BKMM12, Theorem 8.5] and [EMR, Theorem A].

Fundamental groups of non-geometric 3–manifolds are HHS of rank 2, [BHS15]. For these groups, the above theorem allows us to recover the follwoing quasiflats theorem which was first established by Kapovich–Leeb [KL97a]:

Corollary D (Quasiflats theorem for non-geometric 3–manifolds; [KL97a])**.** *Any top-dimensional quasiflat in a non-geometric* 3*–manifold is uniformly close to a finite union of standard flats.*

For CAT(0) cube complexes, we obtain the following result which generalizes the main theorems of [BKS16] and [Hua14b] in certain cocompact cases:

Corollary E (Quasiflats theorem for cubulated groups)**.** *Let $\mathcal{X}$ be a CAT(0) cube complex admitting a* factor system *in the sense of* [BHS17b]*. Let ν be the maximum dimension of an ℓ_1–isometrically embedded cubical orthant in $\mathcal{X}$. Let $f\colon \mathbb{R}^\nu \to \mathcal{X}$ be a quasi-isometric embedding. Then $\mathsf{d}_{haus}(f(\mathbb{R}^\nu), \cup_{i=1}^k Q_i) < \infty$, where each Q_i can be chosen to be either:*

- *an ℓ^1–isometrically embedded copy of the standard cubical tiling of $[0,\infty)^\nu$, or*
- *a $CAT(0)$–isometrically embedded copy of $[0,\infty)^\nu$ with the Euclidean metric.*

It was established in [BHS17b] that all CAT(0) cube complexes with proper, cocompact, cospecial group actions admit factor systems. More generally, it is shown in [HS] that a CAT(0) cube complex $\mathcal{X}$ has a factor system whenever it admits a proper cocompact action by a group G satisfying any one of a number of natural algebraic conditions, e.g., finite height for hyperplane stabilizers or other weak versions of virtual cospecialness of the G–action. In fact, that paper contains a characterization of actions that give rise to a factor system. We are not aware of any proper CAT(0) cube complex that admits a proper cocompact group action but does not contain a factor system.

Proof. As shown in [BHS17b], $\mathcal{X}^{(1)}$ with the combinatorial metric admits an HHS structure based on the construction in [BHS17b, Section 8]. In particular, the hierarchy paths/rays in $\mathcal{X}^{(1)}$ are combinatorial geodesics, so standard ν–orthants (which are products of hierarchy rays) can be taken to be ℓ_1–embedded copies of the standard cubical tiling of $[0,\infty)^\nu$. By Theorem A we are done, if we choose all our Q_i to be of the first type listed above.

To conclude, it suffices to produce N so that for any ℓ_1–isometric embedding $o\colon \prod_{i=1}^{\nu}\gamma_i \to \mathcal{X}$ with γ_i a combinatorial geodesic ray, there is a CAT(0) orthant o' with $\mathsf{d}_{haus}(\mathrm{im}(o),o') \leqslant N$. For each i, let $\mathcal{Y}_i$ be the convex hull of γ_i, i.e., the intersection of all combinatorial halfspaces containing γ_i. Then the hull of $\mathrm{im}(o)$ decomposes as $\prod_{i=1}^{\nu}\mathcal{Y}_i$. Since $\mathcal{Y}_i$ contains a CAT(0)–geodesic ray crossing all hyperplanes, it suffices to show that $\mathcal{Y}_i$ lies uniformly close to γ_i. But if there is no such bound, then for any m, we can choose o so that for some i, we have an ℓ_1–isometric embedding $[0,m]^2 \to \mathcal{Y}_i$, and thus an ℓ_1–isometric embedding $[0,m]^2\times[0,\infty)^{\nu-1} \to \mathcal{X}$. Cocompactness would then allow us to produce a $(\nu+1)$–dimensional cubical orthant in $\mathcal{X}$, which is impossible by our choice of ν. □

Observe that the quasiflats in Corollary E may have dimension strictly less than the dimension of $\mathcal{X}$, since a cube complex may contain cubes of high dimension that are not contained in cubical orthants; for instance, there exists hyperbolic (and hence rank one) cubulated groups, whose associated cube complexes have arbitrarily large dimension. In this sense, this corollary is stronger than the cases covered in [Hua14b], since our result applies even if the dimension is larger than the rank; on the other hand, in practice the construction of a factor systems relies on a geometric group action, which is a hypothesis not needed in the context of [Hua14b].

Approximating with cube complexes. In Section 2, we introduce a new tool for studying hierarchically hyperbolic spaces, which we expect will have a number of applications beyond those of this paper. Roughly, this theorem says that "convex hulls" of finite sets, denoted $H_\theta(A)$, are approximated by finite CAT(0) cube complexes:

Theorem F (Approximation of convex hulls in HHS by CAT(0) cube complexes)**.** *Let $\mathcal{X}$ be an asymphoric HHS of rank ν. Then for any N there exists C so that the following holds. Let $A\subseteq\mathcal{X}$ have cardinality at most N. Then there exists a CAT(0) cube complex $\mathcal{Y}$ of dimension at most ν and a C–quasimedian (C,C)–quasi-isometry $\mathfrak{p}_A\colon \mathcal{Y}\to H_\theta(A)$.*

Any HHS is coarse median in the sense of [Bow13], as shown in [BHS15, Section 7]. However, since Theorem F provides an approximation of the entire convex hull, the "cubical approximations of finite sets" provided by Theorem F have much stronger properties than the "cubical approximations of finite sets" provided by the definition of a coarse median space, or the metric approximation result given in [Zei16, Theorem 6.2]. In fact, the *quasimedian map* from a finite median algebra provided by the coarse median property can be very far from having uniformly (hierarchically) quasiconvex image. To see the distinction, consider the case where $\mathcal{X}=\mathbb{Z}^2$ and $A=\{(0,0),(n,n)\}$ for some $n\geqslant 0$. Then the $\mathcal{Y}$ provided by Theorem F is a n–by–n square, while the 2–point median algebra $\{(0,0),(n,n)\}$ satisfies the requirements of the definition of a coarse median space, and is a "metric approximation" in the sense of [Zei16] when endowed with the natural metric.

Theorem F allows us to control the rank of $\mathcal{X}$ as a coarse median space more precisely than we did in [BHS15]; see Corollary 2.15. This also leads to a characterization of hierarchically hyperbolic spaces which are hyperbolic, Corollary 2.16.

Induced quasi-isometries on factored spaces and quasi-isometric classification. In [BHS17a], we introduced the notion of *factored spaces* of an HHS. These are obtained from a given HHS by "coning off" a collection of product regions, and they are HHS themselves with respect to a substructure of the original HHS. Factored spaces are central in the proof of finite asymptotic dimension [BHS17a], and naturally occurring examples include: the Weil-Petersson metric on Teichmüller space, which is (quasi-isometric to) a factored space of the corresponding mapping class group; and, in any HHS, a space quasi-isometric to the image of $\mathcal{X}$ in $\mathcal{C}S$ for the $\sqsubseteq$–maximal element S (e.g., $\mathcal{C}S$ is the curve graph of S when S is a surface and $\mathcal{X} = \mathcal{MCG}(S)$).

In Theorem 6.2 we use the Quasiflats Theorem as a starting point to show that the image of any quasiflat in a certain factored space is bounded. For now, we just state a new result about mapping class groups which is a special case of Theorem 6.2:

Theorem G (Quasiflats have finite diameter $\mathcal{C}S$ projection)**.** *Let* $(\mathcal{X}, \mathfrak{S})$ *be the mapping class group of a non-sporadic surface* S*. Then for every* K *there exists* L *so that any* (K,K)*–quasi-isometric embedding* $f\colon \mathbb{R}^\nu \to \mathcal{X}$ *satisfies* $\mathsf{diam}_{\mathcal{C}S}(\pi_S(f(\mathbb{R}^\nu))) \leqslant L$*.*

As Corollary 6.3 we prove that any quasi-isometry between HHS satisfying a mild condition induces a quasi-isometry of the factored spaces obtained by coning off the standard product regions containing top-dimensional quasiflats. This is very important because one can extract further information about the original quasi-isometry from the induced quasi-isometry on factored spaces, and even take further factored spaces for additional data. This is totally unexplored territory, since, for example, it provides a way to study quasi-isometries of CAT(0) cube complexes that requires leaving the world of cube complexes.

We expect this strategy to be crucial to prove quasi-isometric rigidity results for, say, right-angled Artin and Coxeter groups. We discuss this in more detail below; for now we just give an example of two right-angled Artin groups whose quasi-isometry classes can be distinguished using this method, but not by any other known methods: see Figure 1. The obstruction to their being quasi-isometric is that, despite having the same rank, their factored spaces as in Corollary 6.3 have different rank (which is a quasi-isometry invariant by Theorem 1.14). We note that the graphs we chose do not fit the hypotheses of [Hua14a, Hua16], or that of any other class of right-angled Artin groups which have been classified including those considered in [BN08, BJN10, BKS08].

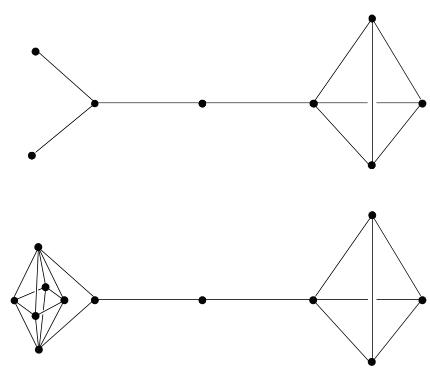

FIGURE 1. The right-angled Artin groups associated to the two graphs both have rank 4. However, the 4-dimensional flats get collapsed in the corresponding factored spaces, leaving only 2-dimensional flats in the case of the first RAAG, while there are 3-dimensional flats that persist in the case of the second RAAG.

Induced automorphisms of combinatorial data and quasi-isometric rigidity. The Quasiflats Theorem provides a powerful tool for proving quasi-isometric rigidity results for classes of HHS, for example right-angled Artin and Coxeter groups. In fact, the set of quasiflats and, more importantly, their intersection patterns, can be easily converted into purely combinatorial data. In good cases, one can extract from the output of the Quasiflats Theorem (and with basically no further knowledge about the geometry of the HHS) an automorphism of a combinatorial structure encoding the data, and therefore reduce proving quasi-isometric rigidity to proving that a certain combinatorial structure is "rigid". The kind of combinatorial structure that the reader should keep in mind is $\mathfrak{S}$ endowed with the partial order given by nesting, $\sqsubseteq$, and the symmetric relation of orthogonality, $\perp$.

Rather than a general but complicated statement, we give a template for this procedure. In Theorem 5.8 we give an example of the combinatorial automorphism one can extract from a quasi-isometry, under additional assumptions on the HHS. These additional assumptions are satisfied by mapping class groups. Accordingly, in Section 5.2, we use Theorem 5.8 to give a short new proof of quasi-isometric rigidity of mapping class groups that relies on much simpler combinatorial considerations than previous proofs, cf. [BKMM12, Bow15, Ham07].

Theorem H (QI rigidity for mapping class groups; [BKMM12])**.** *Let $\mathcal{X}$ be the the mapping class group of a non-sporadic surface S. Then for any K there exists L so that: for each quasi-isometry $f\colon \mathcal{X} \to \mathcal{X}$ there exists a mapping class g so that f L–coarsely coincides with left-multiplication by g.*

Theorem 5.8 applies to other spaces and groups as well, including, for example, right-angled Artin groups with no triangles and no leaves in their presentation graph, and fundamental groups of non-geometric graph manifolds. Variations of Theorem 5.8 can be tailored to treat other families of groups as well.

In the case of mapping class groups, there is no need to pass to factored spaces, but in other contexts (e.g., the right-angled Artin groups in Figure 1) the induced quasi-isometries on factored spaces provide extra combinatorial data.

In the study of right-angled Artin and Coxeter groups our results allow one to reduce the question of quasi-isometric rigidity to the following type of combinatorial problem, which we believe is of independent interest. Let Γ be a finite simplicial graph, and let B_Γ be either the associated right-angled Artin group or the associated right-angled Coxeter group. Recall from [BHS17b, Section 8] that the standard hierarchically hyperbolic structure on such a group is obtained by setting $\mathfrak{S}_\Gamma = \{gB_\Lambda\}/_\sim$, where $g \in B_\Gamma$ and Λ is an induced subgraph of Γ, where $\sim$ is the equivalence relation defined by $gB_\Lambda \sim hB_\Lambda$ if $g^{-1}h \in B_{star(\Lambda)}$, and where $star(\varnothing) = \Gamma$ (i.e. $g^{-1}h$ commutes with each $b \in B_\Lambda$). Declare $[gB_\Lambda] \sqsubseteq [gB_{\Lambda'}]$ if $\Lambda \subseteq \Lambda'$ and $[gB_\Lambda] \perp [gB_{\Lambda'}]$ if $\Lambda \subseteq link(\Lambda')$ and $\Lambda' \subseteq link(\Lambda)$. Answers to the following can be used to obtain results on the problems of quasi-isometric rigidity and classification:

Problem I. *Study the automorphism group* $\mathrm{Aut}(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$ *of* $(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$*. When is every element of* $\mathrm{Aut}(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$ *induced by left multiplication by an element of B_Γ? When is every element of* $\mathrm{Aut}(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$ *"induced" by an automorphism of B_Γ? (Not all automorphisms of B_Γ need to "induce" an automorphism of* $(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$*; which ones do?)*

Theorem 5.8 states that, under three natural assumptions, a quasi-isometry $f\colon (\mathcal{X}, \mathfrak{S}) \to (\mathcal{Y}, \mathfrak{T})$ induces a bijection from the set of *hinges* of $\mathcal{X}$ to that of $\mathcal{Y}$; a *hinge* in $\mathcal{X}$ is a pair (U, p) with $U \in \mathfrak{S}$ and $p \in \partial \mathcal{C}U$, where U has the additional property that $U \in \{U_i\}_{i=1}^{\nu}$ where ν is the rank of $\mathcal{X}$, each $\mathcal{C}U_i$ is unbounded, and the U_i are pairwise-orthogonal.

Since it preserves orthogonality, this bijection determines a simplicial isomorphism from the union of the top-dimensional simplices of the HHS boundary $\partial\mathcal{X}$ to $\partial\mathcal{Y}$ (see [DHS17] for more on the HHS boundary and its simplices). One should be able to articulate natural conditions defining a subclass of HHS for which one can use this map, perhaps in conjunction with Section 6, to pass from a quasi-isometry to a map between HHS boundaries.

Outline. Section 1 contains background on hierarchically hyperbolic spaces, wallspaces/cube complexes, median and coarse median spaces, and asymptotic cones. In Section 2 we build walls in hulls of finite sets, proving Theorem F. The main goal of Section 3 is to prove Corollary 3.8, showing that balls in quasiflats in an HHS can be uniformly well-approximated by hulls of uniformly finite sets of points. In Section 4, we develop background on standard orthants in HHS, and then prove Theorem A, as well as stronger versions in which we control both the number of standard orthants (using a volume growth argument) and the distance from the quasiflat to the approximating orthants, in terms of the quasi-isometry constants. In Section 5, we impose additional assumptions on an HHS enabling one to study the effect of quasi-isometries on the underlying combinatorial structure; see Theorem 5.8; it is in this section that we give a new proof of quasi-isometric rigidity of the mapping class group, i.e., Theorem H. Finally, in Section 6, we discuss factored spaces, proving Theorem 6.2 and its important consequence yielding induced quasi-isometries of factored spaces, Corollary 6.3.

Acknowledgments. We would like to thank an anonymous ancient Greek scholar for early input on this paper, in particular for suggesting the adjective "asymphoric," derived from the word $\sigma\upsilon\mu\phi o\rho\acute{\alpha}$ and roughly meaning "mishap averted."

The authors thank Jingyin Huang, for pointing out that the RAAGs appearing in Figure 1 in an earlier version of this paper could be distinguished by using [Hua14a, Theorem 3.28] and a clever trick. The authors also thank Brian Bowditch for several helpful comments.

The authors were supported by the National Science Foundation under Grant No. DMS-1440140 at the Mathematical Sciences Research Institute in Berkeley during Fall 2016 program in Geometric Group Theory. Hagen was supported by EPSRC grant EP/L026481/1. We thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the program *Non-positive curvature group actions and cohomology* where work on this paper was done. This work was partly supported by EPSRC grant no. EP/K032208/1. In Section 2, we build walls; the Newton Institute paid for them.

1. Background

1.1. **Hierarchically hyperbolic spaces.** Throughout this paper, we work with a *hierarchically hyperbolic space*, which is a pair $(\mathcal{X}, \mathfrak{S})$ with some additional extra structure described in Definition 1.1 of [BHS15]. Roughly, an HHS consists of:

- a quasigeodesic metric space $\mathcal{X}$;
- a set of uniformly hyperbolic spaces $\{\mathcal{C}U : U \in \mathfrak{S}\}$;
- uniformly coarsely-Lipschitz coarsely-surjective maps $\pi_U \colon \mathcal{X} \to \mathcal{C}U$;
- three relations $\sqsubseteq$ (a partial order), $\perp$ (an anti-reflexive symmetric relation), $\pitchfork$ (the complement of $\sqsubseteq$ and $\perp$) on $\mathfrak{S}$;
- a unique $\sqsubseteq$–maximal element of $\mathfrak{S}$, and a uniform bound on the length of $\sqsubseteq$–chains in $\mathfrak{S}$;
- for $U \sqsubsetneq V$ or $U \pitchfork V$, a uniformly bounded set ρ^U_V;
- for $U \sqsubsetneq V$, a coarse map $\rho^V_U : \mathcal{C}V \to \mathcal{C}U$.

Definition 1.1 of [BHS15] consists of several axioms governing this data. The definition and basic properties of HHS were first laid out in [BHS17b]; below we list [BHS15] as the primary reference since a few of the properties were first established there and this provides for unified notation. The properties of HHS which are central to this article are listed below.

The first one says that the "coordinates" $(\pi_U(x))_{U \in \mathfrak{S}}$ for some $x \in \mathcal{X}$ cannot be arbitrary. In fact, for certain pairs U, V there are conditions that need to be satisfied by $\pi_U(x), \pi_V(x)$.

There is no condition for $U \perp V$, which corresponds to the fact that in this case U, V should be thought of as factors of a product region, as we will see later.

Axiom 1.1 (Consistency axioms)**.** *Let $(\mathcal{X}, \mathfrak{S})$ be hierarchically hyperbolic. Then there is a constant $E = E(\mathcal{X}, \mathfrak{S})$ so that the following hold for all $x \in \mathcal{X}$ and $U, V, W \in \mathfrak{S}$:*

- *if $V \pitchfork W$, then*
$$\min\left\{\mathsf{d}_W(\pi_W(x), \rho^V_W), \mathsf{d}_V(\pi_V(x), \rho^W_V)\right\} \leqslant E;$$
- *if $V \subsetneq W$, then*
$$\min\left\{\mathsf{d}_W(\pi_W(x), \rho^V_W), \mathsf{diam}_{\mathcal{C}V}(\pi_V(x) \cup \rho^W_V(\pi_W(x)))\right\} \leqslant E.$$

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

The following theorem says that we can compute distances in $\mathcal{X}$ in terms of distances in the various $\mathcal{C}U$, thereby reducing the study of the geometry of $\mathcal{X}$ to that of the family of hyperbolic spaces $\{\mathcal{C}U\}_{U \in \mathfrak{S}}$. Notice that a special case of the distance formula is that, roughly speaking, if $x, y \in \mathcal{X}$ are so that $\pi_U(x), \pi_U(y)$ are close for each U, then x, y are close in $\mathcal{X}$ (this is the uniqueness axiom).

We write $A \asymp_{K,C} B$ if $A/K - C \leqslant B \leqslant KA + C$. Also, we let $\{\!\{A\}\!\}_s = A$ if $A \geqslant s$, and $\{\!\{A\}\!\}_s = 0$ otherwise. Moreover, we denote $\mathsf{d}_W(x, y) = \mathsf{d}_{\mathcal{C}W}(\pi_W(x), \pi_W(y))$ (the distance between x and y from the point of view of W).

Theorem 1.2 (Distance Formula; [BHS15])**.** *Let $(X, \mathfrak{S})$ be hierarchically hyperbolic. Then there exists s_0 such that for all $s \geqslant s_0$ there exist constants K, C such that for all $x, y \in \mathcal{X}$,*
$$\mathsf{d}_{\mathcal{X}}(x, y) \asymp_{K,C} \sum_{W \in \mathfrak{S}} \{\!\{\mathsf{d}_W(x, y)\}\!\}_s .$$

Pairs of points in HHS are connected by special quasi-geodesics, called hierarchy paths:

Theorem 1.3 (Existence of Hierarchy Paths; [BHS15])**.** *Let $(\mathcal{X}, \mathfrak{S})$ be hierarchically hyperbolic. Then there exists D so that any $x, y \in \mathcal{X}$ are joined by a D-hierarchy path, i.e. a (D, D)–quasi-geodesic projecting to an unparameterized (D, D)–quasi-geodesic in $\mathcal{C}U$ for each $U \in \mathfrak{S}$.*

The following theorem says that the conditions in the consistency axiom in fact characterize the coordinates that are (coarsely) realized by a point in $\mathcal{X}$.

Theorem 1.4 (Realization of consistent tuples; [BHS15])**.** *For each $\kappa \geqslant 1$ there exist $\theta_e, \theta_u \geqslant 0$ such that the following holds. Let $\vec{b} \in \prod_{W \in \mathfrak{S}} 2^{\mathcal{C}W}$ be κ–consistent (*[BHS15, Definition 1.16]*); for each W, let b_W denote the $\mathcal{C}W$–coordinate of $\vec{b}$.*

Then there exists $x \in \mathcal{X}$ so that $\mathsf{d}_W(b_W, \pi_W(x)) \leqslant \theta_e$ for all $\mathcal{C}W \in \mathfrak{S}$. Moreover, x is coarsely unique *in the sense that the set of all x which satisfy $\mathsf{d}_W(b_W, \pi_W(x)) \leqslant \theta_e$ in each $\mathcal{C}W \in \mathfrak{S}$, has diameter at most θ_u.*

The following says that when moving along a hierarchy path γ, in order to change projection to $\mathcal{C}U$, when $U \subsetneq V$, one must pass close in $\mathcal{C}V$ to a specific point, namely ρ^U_V.

Lemma 1.5. *(Bounded geodesic image) Let $\mathcal{X}$ be a hierarchically hyperbolic space. There exists B so that the following holds. Let $W \in \mathfrak{S}$, $V \subsetneq W$. Suppose that γ is a geodesic in $\mathcal{C}W$ with $\gamma \cap \mathcal{N}_B(\rho^V_W)) = \emptyset$. Then $\mathsf{diam}_{\mathcal{C}V}(\rho^W_V(\gamma)) \leqslant B$.*

Moreover, suppose $x, y \in \mathcal{X}$ and that there exists a geodesic γ in $\mathcal{C}W$ from $\pi_W(x)$ to $\pi_W(y)$ so that $\gamma \cap \mathcal{N}_B(\rho^V_W)) = \emptyset$. Then $\mathsf{d}_V(x, y) \leqslant B$.

The following is a variation of [BHS15, Lemma 2.5]. For $V \in \mathfrak{S}$, we denote $\mathfrak{S}_V = \{U \in \mathfrak{S} : U \sqsubseteq V\}$.

Lemma 1.6 (Passing large projections up the $\sqsubseteq$–lattice)**.** *There exists E with the following property. For every $C \geqslant 0$ there exists $N_0 = N_0(C)$ with the following property. Let $V \in \mathfrak{S}$, let $x, y \in \mathcal{X}$, and let $\{V_i\}_{i=1}^{N_0} \subseteq \mathfrak{S}_V$ be distinct and satisfy $\mathsf{d}_{V_i}(x, y) \geqslant E$. Then there exists $W \in \mathfrak{S}_V$ and i, j so that $V_i, V_j \subsetneq W$ and $\mathsf{d}_W(\rho_W^{V_i}, \rho_W^{V_j}) \geqslant C$.*

Proof. First of all, we choose constants. Let $B \geqslant 1$ be the constant from Lemma 1.5, and suppose that B is also an upper bound on the diameter of ρ_V^U for any $U \subsetneq V$. Moreover, supposed $B \geqslant D$, for D as in Theorem 1.3, and moreover that (D, D)–quasi-geodesics in a δ–hyperbolic space stay B–close to geodesics with the same endpoints, where δ is a hyperbolicity constant for all the $\mathcal{C}U$.

If $U \in \mathfrak{S}$ is $\sqsubseteq$–minimal, we say that its *level* is 1. Inductively, $U \in \mathfrak{S}$ has level k if it is $\sqsubseteq$–minimal among all $V \in \mathfrak{S}$ not of level $\leqslant k - 1$. The proof is by induction on the level k of a $\sqsubseteq$-minimal $V \in \mathfrak{S}$ into which each V_i is nested, with $E = 100kB$. The base case $k = 1$ is empty. Suppose that the statement holds for a given $N = N(k)$ when the level of V as above is at most k. Suppose instead that $|\{V_i\}| \geqslant N(k+1)$ (where $N(k+1)$ is a constant much larger than $N(k)$ that will be determined shortly) and there exists a $\sqsubseteq$-minimal $V \in \mathfrak{S}$ of level $k + 1$ into which each V_i is nested. There are two cases.

If $\max_{i,j}\{\mathsf{d}_V(\rho_V^{V_i}, \rho_V^{V_j})\} \geqslant C$, then we are done. Hence, suppose not. All the $\rho_V^{V_i}$ lie B–close to a geodesic $[\pi_V(x), \pi_V(y)]$ by bounded geodesic image, and by the assumption they all lie close to a sub-geodesic of length $C + 10B$. Hence, we can replace x, y with suitable x', y' on a hierarchy path from x to y chosen so that

- $\mathsf{d}_V(x', y') \leqslant C + 100B$,
- $\pi_V(x'), \pi_V(y')$ lie B–close to a geodesic $[\pi_V(x), \pi_V(y)]$, and
- the geodesics $[\pi_V(x), \pi_V(x')]$, $[\pi_V(y), \pi_V(y')]$ do not pass B–close to any $\rho_V^{V_i}$.

By Lemma 1.5, $\mathsf{d}_{V_i}(x', y') \geqslant 100kB$, since $\mathsf{d}_{V_i}(x', y')$ is approximately equal to $\mathsf{d}_{V_i}(x, y)$.

The large link axiom ([BHS15, Definition 1.1.(6)]) implies that there exists $K = K(C + 100B)$ and $T_1, \ldots, T_K$, each properly nested in V (thus of level strictly less than $k + 1$), so that any V_i is nested in some T_j. In particular, if $N(k+1) \geqslant KN(k)$, there exists j so that $\geqslant N(k)$ elements of $\{V_i\}$ are nested into T_j. By the induction hypothesis, we are done. $\square$

Notation 1.7. In the remainder of the paper, following [BHS15, Remark 1.5], we fix a constant E larger than each of the constants in [BHS15, Definition 1.1] and also satisfying the conclusion of Lemma 1.6.

Definition 1.8 (Relevant)**.** Given points $x, y \in \mathcal{X}$, we say that $U \in \mathfrak{S}$ is *relevant* (with respect to x, y and a constant $\theta > 0$) if $\mathsf{d}_U(x, y) > \theta$. Denote by $\mathbf{Rel}_\theta(x, y)$ the set of relevant elements.

Definition 1.9 (Rank)**.** The *rank* $\nu = \nu(\mathcal{X}, \mathfrak{S})$ of the HHS $(\mathcal{X}, \mathfrak{S})$ is the maximal n so that there exist pairwise orthogonal $U_1, \ldots, U_n \in \mathfrak{S}$ for which $\pi_{U_i}(\mathcal{X})$ is unbounded for all i.

Standard product regions are a standard useful tool; see [BHS17b, Section 13] and [BHS15]. These products are built out of the following two spaces, which we define abstractly, but often implicitly identify with their images as subsets of $\mathcal{X}$.

Definition 1.10. Recall that $\mathfrak{S}_U = \{V \in \mathfrak{S} \mid V \sqsubseteq U\}$. Fix $\kappa \geqslant E$ and let F_U be the set of κ–consistent tuples in $\prod_{V \in \mathfrak{S}_U} 2^{\mathcal{C}V}$.

Definition 1.11. Let $\mathfrak{S}_U^\perp = \{V \in \mathfrak{S} \mid V \perp U\}$. Fix $\kappa \geqslant E$ and let E_U be the set of κ–consistent tuples in $\prod_{V \in \mathfrak{S}_U^\perp} 2^{\mathcal{C}V}$.

Definition 1.12 (Standard product regions in $\mathcal{X}$)**.** Given $\mathcal{X}$ and $U \in \mathfrak{S}$, there are coarsely well-defined maps $\phi^{\sqsubseteq}, \phi^\perp \colon F_U, E_U \to \mathcal{X}$ which extend to a coarsely well-defined map $\phi_U \colon F_U \times$

$E_U \to \mathcal{X}$. Indeed, for each $(\vec{a}, \vec{b}) \in F_U \times E_U$, and each $V \in \mathfrak{S}$, the coordinate $(\phi_U(\vec{a}, \vec{b}))_V$ is defined as follows. If $V \sqsubseteq U$, then $(\phi_U(\vec{a}, \vec{b}))_V = a_V$. If $V \perp U$, then $(\phi_U(\vec{a}, \vec{b}))_V = b_V$. If $V \pitchfork U$, then $(\phi_U(\vec{a}, \vec{b}))_V = \rho^U_V$. Finally, if $U \subsetneq V$, let $(\phi_U(\vec{a}, \vec{b}))_V = \rho^U_V$. We refer to $F_U \times E_U$ as a *standard product region.*

1.1.1. *Rank as a quasi-isometry invariant.* We now introduce a technical assumption on the HHS that we will assume throughout the paper. This condition is satisfied by all HHG; it is also satisfied for all naturally occurring examples of HHS. We impose it in order to rule out product regions with bounded but arbitrarily large factors. This hypothesis plays an important role in bounding the dimension of the CAT(0) cube complexes approximating hulls of finitely many point, and our theorems fail to hold without this assumption. Nonetheless, our results likely have analogues that hold in the absence of this hypothesis, but would require custom-tailoring to the situation at hand.

Definition 1.13 (Asymphoric)**.** We say that the HHS $(\mathcal{X}, \mathfrak{S})$ of rank ν is *asymphoric* if there exists a constant C with the property that there does not exist a set of $\nu + 1$ pairwise orthogonal elements U of $\mathfrak{S}$ where each $\mathcal{C}U$ has diameter at least C. In this case, without loss of generality, we assume that E is chosen to be at least as large as C.

For completeness, we remark that a result from [BHS17b] implies that the rank is a quasi-isometry invariant of asymphoric HHS:

Theorem 1.14 (Quasi-isometry invariance of rank)**.** *Let $(\mathcal{X}, \mathfrak{S})$ be an asymphoric HHS. Then the rank ν of $\mathcal{X}$ coincides with the maximal n for which there exists K and (K, K)–quasi-isometric embeddings $f\colon (B_R(0) \subseteq \mathbb{R}^n) \to \mathcal{X}$ for all $R \geqslant 0$. In particular, the rank is a quasi-isometry invariant of asymphoric HHS.*

Proof. It is easy to construct a quasi-isometric embeddings of balls in $\mathbb{R}^n$ starting from n pairwise orthogonal elements U of $\mathfrak{S}$ with unbounded $\pi_U(\mathcal{X})$. Hence, we have to show that if there exist quasi-isometric embeddings as in the statement, then n is at most the rank. This is because, by [BHS17b, Theorem 13.11.(2)], there exists an asymptotic cone $\boldsymbol{\mathcal{X}}$ where a copy of the unit ball in $\mathbb{R}^n$ is contained in an ultralimit of standard boxes. These are products of intervals contained in a subspace decomposing as product whose factors are various subspaces F_U, so that any ultralimit of standard boxes in $\boldsymbol{\mathcal{X}}$ is homeomorphic to a subset of $\mathbb{R}^\nu$ because $\mathcal{X}$ is asymphoric. Hence, $n \leqslant \nu$, as required. □

1.2. **Hulls and gates.** Sets in an HHS have *hulls*, built from convex hulls in hyperbolic spaces:

Definition 1.15 (Hull of a set; [BHS15])**.** For each $A \subset \mathcal{X}$ and $\theta \geqslant 0$, let the *hull*, $H_\theta(A)$, be the set of all $p \in \mathcal{X}$ so that, for each $W \in \mathfrak{S}$, the set $\pi_W(p)$ lies at distance at most θ from $\mathrm{hull}_{\mathcal{C}W}(A)$, the convex hull of A in the hyperbolic space $\mathcal{C}W$ (that is to say, the union of all geodesics in $\mathcal{C}W$ joining points of A). Note that $A \subset H_\theta(A)$.

Definition 1.16 (Hierarchical quasiconvexity [BHS15])**.** Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then $\mathcal{Y} \subseteq \mathcal{X}$ is *k–hierarchically quasiconvex*, for some $k\colon [0, \infty) \to [0, \infty)$, if the following hold:

(1) For all $U \in \mathfrak{S}$, the projection $\pi_U(\mathcal{Y})$ is a $k(0)$–quasiconvex subspace of the δ–hyperbolic space $\mathcal{C}U$.
(2) For all $\kappa \geqslant 0$ and κ-consistent tuples $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{\mathcal{C}U}$ with $b_U \subseteq \pi_U(\mathcal{Y})$ for all $U \in \mathfrak{S}$, each point $x \in \mathcal{X}$ for which $\mathsf{d}_U(\pi_U(x), b_U) \leqslant \theta_e(\kappa)$ (where $\theta_e(\kappa)$ is as in Theorem 1.4) satisfies $\mathsf{d}(x, \mathcal{Y}) \leqslant k(\kappa)$.

Proposition 1.17. [BHS15, Lemma 6.2] *There exists θ_0 so that for each $\theta \geqslant \theta_0$ there exists $\kappa\colon \mathbb{R}_+ \to \mathbb{R}_+$ so that for each $A \subset \mathcal{X}$ the set $H_\theta(A)$ is κ–hierarchically quasiconvex.*

Remark 1.18. We fix once and for all $\theta \geqslant \theta_0$.

We now recall a construction from Section 5 of [BHS15], namely the *gate map* to a hierarchically quasiconvex subspace, and prove some additional facts about it. We fix a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$.

Let $A \subset \mathcal{X}$ be κ–hierarchically quasiconvex. Recall, this implies that for each $U \in \mathfrak{S}$, the set $\pi_U(A)$ is $\kappa(0)$–quasiconvex in $\mathcal{C}U$ and there is thus a coarse closest-point projection $p_{U,A}\colon \mathcal{C}U \to \pi_U(A)$. Define a *gate map* $\mathfrak{g}_A\colon \mathcal{X} \to A$ as follows: given $x \in \mathcal{X}$, for each $U \in \mathfrak{S}$ let $b_U = p_{U,A}(x)$. In [BHS15, Section 5] we show that the tuple $(b_U)_{U\in\mathfrak{S}}$ is uniformly (depending on $\kappa(0)$) consistent, so Theorem 1.4 and hierarchical quasiconvexity of A produce a coarsely unique point $\mathfrak{g}_A(x) \in A$ such that $\pi_U(\mathfrak{g}_A(x))$ uniformly coarsely coincides with b_U for all $U \in \mathfrak{S}$.

The following lemma contains a lot of information about the gates of a hierarchically quasiconvex sets A, B. It essentially describes a "bridge" of the form $\mathfrak{g}_A(B) \times H_\theta(A,B)$, for suitable $a \in A, b \in B$ that connects the two. An efficient way to go from $a' \in A$ to $b' \in B$ is to start at a', get to the bridge, cross it, and then go to b'.

The lemma collects more information than we will need in this paper, for future reference. The proof can be safely skipped on first reading.

Lemma 1.19. *For every κ there exists κ', K such that for any κ–hierarchically quasiconvex sets A, B, the following hold.*

(1) $\mathfrak{g}_A(B)$ is κ'–hierarchically quasi-convex.
(2) The composition $\mathfrak{g}_A \circ \mathfrak{g}_B|_{\mathfrak{g}_A(B)}$ is bounded distance from the identity $\mathfrak{g}_A(B) \to \mathfrak{g}_A(B)$.
(3) For any $a \in \mathfrak{g}_A(B), b = \mathfrak{g}_B(a)$, we have a quasi-isometric embedding $f\colon \mathfrak{g}_A(B) \times H_\theta(\{a,b\}) \to \mathcal{X}$ with image $H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$, so that $f(\mathfrak{g}_A(B) \times \{b\})$ K–coarsely coincides with $\mathfrak{g}_B(A)$.

Let $\mathcal{H} = \{U \in \mathfrak{S} : \mathsf{diam}(\mathfrak{g}_A(B)) > K\}$.

(4) For each $p, q \in \mathfrak{g}_A(B)$ and $t \in H_\theta(\{a,b\})$, we have $\mathbf{Rel}_K(f(p,t), f(q,t)) \subseteq \mathcal{H}$.
(5) For each $p \in \mathfrak{g}_A(B)$ and $t_1, t_2 \in H_\theta(\{a,b\})$, we have $\mathbf{Rel}_K(f(p,t_1), f(p,t_2)) \subseteq \mathcal{H}^\perp$.
(6) For each $p \in A, q \in B$ we have

$$\mathsf{d}(p,q) \asymp_{K,K} \mathsf{d}(p, \mathfrak{g}_A(B)) + \mathsf{d}(q, \mathfrak{g}_B(A)) + \mathsf{d}(A,B) + \mathsf{d}(\mathfrak{g}_{\mathfrak{g}_B(A)}(p), \mathfrak{g}_{\mathfrak{g}_B(A)}(q)).$$

Proof. We start with a definition and an observation.

The sets $\mathcal{V}, \mathcal{H}$: Let $\mathcal{V}$ be the set of $V \in \mathfrak{S}$ with $\mathsf{d}_V(A,B) \geqslant 100E\kappa(0)$. As in the statement of the lemma, we define $\mathcal{H}$ to be the set of $H \in \mathfrak{S}$ with $\mathsf{d}_H(a,a') > 10E\kappa(0)$ for some $a, a' \in \mathfrak{g}_A(B)$, say $a = \mathfrak{g}_A(b), a' = \mathfrak{g}_A(b')$ for some $b, b' \in B$. We have $V \perp H$ for all $V \in \mathcal{V}$ and $H \in \mathcal{H}$, by Lemma 1.25 together with the following claim, which can be proved using standard quadrilateral arguments.

Claim 1.20. *$\pi_V(\mathfrak{g}_A(B))$ and $\pi_V(\mathfrak{g}_B(A))$ have diameter $\leqslant 10E\kappa(0)$ for $V \in \mathcal{V}$.*
For $U \in \mathfrak{S} - \mathcal{V}$ and $x \in \mathfrak{g}_A(B)$, $\mathsf{d}_U(x, \mathfrak{g}_B(x)) \leqslant 10E\kappa(0)$.

Assertion (1) **and Assertion** (2)**:** First we claim that $\pi_U(\mathfrak{g}_A(B))$ is uniformly quasiconvex for all $U \in \mathfrak{S}$. Observe that $\pi_U(\mathfrak{g}_A(B))$ uniformly coarsely coincides with $p_{U,A}(\pi_U(B))$. On the other hand, (uniform) quasiconvexity of $\pi_U(B)$ and a thin quadrilateral argument show that $p_{U,A}(\pi_U(B))$ is uniformly quasiconvex, as required.

We now verify that $\mathfrak{g}_A(B)$ satisfies the second part of the definition of hierarchical quasiconvexity. To that end, let $(t_U)_{U\in\mathfrak{S}}$ be a consistent tuple so that $t_U = p_{U,A}(b_U)$ for some $b_U \in \pi_U(B)$ for each $U \in \mathfrak{S}$. Theorem 1.4 and hierarchical quasiconvexity of A provide a realization point $x \in A$ for (t_U).

To complete the proof of hierarchical quasiconvexity, we must show that in fact x lies uniformly close to $\mathfrak{g}_A(B)$. Let $y = \mathfrak{g}_A(\mathfrak{g}_B(x))$. Since $y \in \mathfrak{g}_A(B)$, it suffices to show that x

and y are uniformly close. To do so, we show that $\pi_U(x), \pi_U(y)$ are uniformly close for each $U \in \mathfrak{S}$, but this follows by considering the two possibilities for U covered by Claim 1.20. This proves Assertion (1).

For $b \in B$, Claim 1.20 can be applied as above to show that $\pi_U(\mathfrak{g}_A(\mathfrak{g}_B(\mathfrak{g}_A(b))))$ uniformly coarsely coincides with $\pi_U(\mathfrak{g}_A(b))$ for each $U \in \mathfrak{S}$, and hence $\mathfrak{g}_A(\mathfrak{g}_B(\mathfrak{g}_A(b)))$ uniformly coarsely coincides with $\mathfrak{g}_A(b)$ for all $b \in B$, thus proving Assertion (2).

Defining f: Fix $a \in \mathfrak{g}_A(B)$. Choose $b'' \in B$ so that $a = \mathfrak{g}_A(b'')$, and let $b = \mathfrak{g}_B(a)$. Note that $100E\kappa(0) \leqslant \mathsf{d}_V(a,b) \leqslant \mathsf{d}_V(A,B) + 20E\kappa(0)$ for $V \in \mathcal{V}$; the second inequality here follows from Claim 1.20. Since $a \in A$ and $b \in B$ we also have $\mathsf{d}_V(A,B) \leqslant \mathsf{d}_V(a,b)$. For each fixed $a' \in \mathfrak{g}_A(B)$ (up to bounded distance, $a' = \mathfrak{g}_A(b')$ for some $b' \in \mathfrak{g}_B(A)$, by Assertion (2)) and each $U \in \mathfrak{S} - \mathcal{V}$, we set $b_U = \pi_U(a')$. For each $V \in \mathcal{V}$, let γ_V be a geodesic from $\pi_V(a)$ to $\pi_V(b)$ and, for a fixed $h \in H_\theta(\{a,b\})$, set $b_V = \pi_V(h)$, which lies θ–close to γ_V.

Claim 1.21. *Associated to each a', h as above: $(b_W)_{W \in \mathfrak{S}}$ is a uniformly consistent tuple.*

Proof of Claim 1.21. If $W, W' \in \mathfrak{S} - \mathcal{V}$, or if $W, W' \in \mathcal{V}$, then $b_W, b_{W'}$ satisfy any consistency inequality involving W, W', since $b_W, b_{W'}$ coincide with the projections to $\mathcal{C}W, \mathcal{C}W'$ of a common point in those cases.

If $W \in \mathfrak{S} - \mathcal{V}$ and $V \in \mathcal{V}$, then either $W \in \mathcal{H}$ or: $\mathsf{diam}_W(\pi_W(\mathfrak{g}_A(B))) \leqslant 10E\kappa(0)$ and $\mathsf{d}_W(a,b) \leqslant 100E\kappa(0)$. In the first case, $V \perp W$, so there is no consistency inequality to check.

In the second case, if $W \not\perp V$, then a $200E\kappa(0)$–consistency inequality holds, as we now show. Indeed, if $W \pitchfork V$, then $\pi_W(a'), \pi_W(b')$ coarsely coincide, as do $\pi_V(a), \pi_V(a')$ and $\pi_V(b), \pi_V(b')$. At least one of $\pi_V(a')$ or $\pi_V(b')$ is E–far from ρ^W_V, so either $\pi_W(a')$ or $\pi_W(b')$ is uniformly close to ρ^V_W, but these two points coarsely coincide, so $\pi_W(a') = b_W$ is uniformly close to ρ^V_W. The nested cases are similar. □

Assertion (3): Given the consistent tuple provided by Claim 1.21, the realization theorem, Theorem 1.4, then provides a coarsely unique $x \in \mathcal{X}$ realizing (b_W), and we let $f(a', h) = x$. This gives a map $f : \mathfrak{g}_A(B) \times H_\theta(a,b) \to \mathcal{X}$, and one can see using the distance formula that there exists $K = K(\kappa, E)$ so that f is a (K,K)–quasi-isometric embedding. In the next claims, we check that f satisfies the remaining properties of Assertion (3).

Claim 1.22. *$f(\mathfrak{g}_A(B) \times H_\theta(\{a,b\}))$ is coarsely contained in $H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$.*

Proof of Claim 1.22. Let $h \in H_\theta(\{a,b\})$. Let $b' \in B$ and let $x = f(\mathfrak{g}_A(b'), h)$. Let $U \in \mathfrak{S}$. If $U \in \mathcal{V}$, then $\pi_U(x)$ uniformly coarsely coincides with $\pi_U(h)$, which in turn lies θ–close to γ_U by definition. If $U \in \mathfrak{S} - \mathcal{V}$, then $\pi_U(x)$ lies uniformly close to $\pi_U(\mathfrak{g}_A(b'))$. In either case, $\pi_U(x)$ lies uniformly close to a geodesic starting and ending in $\pi_U(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$, so x lies uniformly close to $H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$. □

Claim 1.23. *$H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$ is coarsely contained in the image of f.*

Proof of Claim 1.23. Suppose that $x \in H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$. Let $y = f(\mathfrak{g}_{\mathfrak{g}_A(B)}(x), \mathfrak{g}_{H_\theta(\{a,b\})}(x))$. We claim that $\pi_U(y)$ coarsely coincides with $\pi_U(x)$ for all $U \in \mathfrak{S}$, and hence x coarsely coincides with y. Indeed, suppose that $U \in \mathcal{V}$. By Claim 1.20, we have that $\pi_U(\mathfrak{g}_A(B)), \pi_U(\mathfrak{g}_B(A))$ are uniformly bounded; thus $\pi_U(H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A)))$ coarsely coincides with $\pi_U(H_\theta(\{a,b\}))$. Hence, since $x \in H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$, we have $\pi_U(x)$ coarsely coincides with $\pi_U(\mathfrak{g}_{H_\theta(\{a,b\})}(x))$. By definition, this coarsely coincides with $\pi_U(y)$.

Suppose that $U \in \mathfrak{S} - \mathcal{V}$. Then $\pi_U(\mathfrak{g}_A(B))$ coarsely coincides with $\pi_U(\mathfrak{g}_B(A))$ and hence $\pi_U(H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A)))$ coarsely coincides with $\pi_U(\mathfrak{g}_A(B))$. Hence, since $x \in H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$, we have $\pi_U(x)$ coarsely coincides with $\pi_U(\mathfrak{g}_{\mathfrak{g}_A(B)}(x))$, which coarsely coincides with $\pi_U(y)$ by definition. □

Claim 1.24. *$\mathfrak{g}_B(A)$ coarsely coincides with $f(\mathfrak{g}_A(B) \times \{b\})$.*

Proof of Claim 1.24. By Claim 1.23, $\mathfrak{g}_B(A)$ is coarsely contained in the image of f. Moreover, if $x \in \mathfrak{g}_B(A)$, then $\pi_V(x)$ coarsely coincides with $\pi_V(b)$ for all $V \in \mathcal{V}$, since $b \in \mathfrak{g}_B(A)$ and $\pi_V(\mathfrak{g}_B(A))$ is bounded by Claim 1.20. Hence $\mathfrak{g}_B(A)$ is coarsely contained in $f(\mathfrak{g}_A(B) \times \{b\})$.

Conversely, for any $a' \in \mathfrak{g}_A(B)$, $f(a', b)$ coarsely coincides with $\mathfrak{g}_B(a')$. Indeed, for $V \in \mathcal{V}$, $\pi_V(f(a', b))$ coarsely coincides with $\pi_V(b)$ by definition. But $\pi_V(b) \in \pi_V(\mathfrak{g}_B(A))$, by the choice of b. Since $\pi_V(\mathfrak{g}_B(A))$ is uniformly bounded, $\pi_V(\mathfrak{g}_B(a'))$ coarsely coincides with $\pi_V(b)$ and hence $\pi_V(f(a', b))$.

Let $H \in \mathfrak{S} - \mathcal{V}$. Since $\mathsf{d}_H(A, B) \leqslant 100E\kappa(0)$, we have that $\pi_V(\mathfrak{g}_B(a'))$ coarsely coincides with $\pi_V(a')$. By definition $\pi_V(f(a', b))$ coarsely coincides with $\pi_V(a')$. Hence $f(\mathfrak{g}_A(B) \times \{b\})$ is coarsely contained in $\mathfrak{g}_B(A)$. □

Assertions (4)**,**(5)**:** Let $p, q \in \mathfrak{g}_A(B)$ and $t_1, t_2 \in H_\theta(\{a, b\})$. For sufficiently large K, if $H \in \mathbf{Rel}_K(f(p, t_1), f(q, t_1))$, then by definition $H \in \mathcal{H}$. If $V \in \mathbf{Rel}_K(f(p, t_1), f(p, t_2))$, then by definition $V \in \mathcal{V}$, so $V \in \mathcal{H}^\perp$ by Lemma 1.25, as explained above.

Assertion (6)**:** Let $F = H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$, and consider $p \in A$ and $q \in B$. Assertion (3) and Lemma 1.26 provides K so that

$$\mathsf{d}(\mathfrak{g}_F(p), \mathfrak{g}_F(q)) \asymp_{K,K} \mathsf{d}(A, B) + \mathsf{d}(\mathfrak{g}_{\mathfrak{g}_B(A)}(p), \mathfrak{g}_{\mathfrak{g}_B(A)}(q)),$$

so it suffices to compare $\mathsf{d}(p, q)$ with $\mathsf{d}(p, \mathfrak{g}_F(p)) + \mathsf{d}(\mathfrak{g}_F(p), \mathfrak{g}_F(q)) + \mathsf{d}(q, \mathfrak{g}_F(q))$. The upper bound is just the triangle inequality. For $U \in \mathfrak{S}$, examining a thin quadrilateral shows

$$\begin{aligned} \mathsf{d}_U(p, q) &\geqslant \mathsf{d}_U(p, p_{U,F}(\pi_U(p))) + \mathsf{d}_U(p_{U,F}(\pi_U(p)), p_{U,F}(\pi_U(q))) + \mathsf{d}_U(q, p_{U,F}(\pi_U(q))) - T \\ &\geqslant \mathsf{d}_U(p, \mathfrak{g}_F(p)) + \mathsf{d}_U(\mathfrak{g}_F(p), \mathfrak{g}_F(q)) + \mathsf{d}_U(q, \mathfrak{g}_F(q)) - 10T \end{aligned}$$

for some uniform T. Given $L \geqslant 0$, let $\sigma_L(p, q) = \sum_{U \in \mathfrak{S}} \{\!\{\mathsf{d}_U(p, q)\}\!\}_L$.

By the distance formula (Theorem 1.2), $\mathsf{d}(p, q) \geqslant K_1^{-1}\sigma_{10T}(p, q) - K_1$ for some K_1. Since, $10\sigma_{10T}(p, q) \geqslant \sigma_{100T}(p, \mathfrak{g}_F(p)) + \sigma_{100T}(\mathfrak{g}_F(p), \mathfrak{g}_F(q)) + \sigma_{100T}(\mathfrak{g}_F(p), q)$, the claim follows from another use of the distance formula (on the right, with threshold $100T$). □

Lemma 1.25. *Let $C \geqslant E$ and let $a, b, a', b' \in \mathcal{X}$ and suppose that $H, V \in \mathfrak{S}$ satisfy*

- $\mathsf{d}_V(a, a'), \mathsf{d}_V(b, b') \leqslant C$*;*
- $\mathsf{d}_V(a, b) > 10C$*;*
- $\mathsf{d}_H(a, b), \mathsf{d}_H(a', b') \leqslant C$*;*
- $\mathsf{d}_H(a, a') > 10C$*;*

Then $H \perp V$.

Proof. Suppose $V \pitchfork H$. If $\mathsf{d}_V(a, \rho^H_V) \leqslant E$, then $\mathsf{d}_V(\rho^H_V, b) > 8C$, and hence $\mathsf{d}_V(\rho^H_V, b') > 6C$. Then, by consistency ρ^V_H lies E–close to both $\pi_H(b), \pi_H(b')$, which is impossible since $d_H(b, b') > 6C$. If $\mathsf{d}_V(a, \rho^H_V) > E$, then by consistency $\mathsf{d}_H(a, \rho^V_H) \leqslant E$. Hence $\mathsf{d}_H(a', \rho^V_H) \geqslant 5E$, so by consistency, $\mathsf{d}_V(a', \rho^H_V) \leqslant E$, and we argue as above with a' replacing a.

Suppose $V \subsetneq H$. Since $\mathsf{d}_H(a, a') > 10C$ and $\mathsf{d}_H(b, b') > 6C$, at least one of the pairs a, b or a', b' has the property that geodesics in $\mathcal{C}H$ connecting the corresponding projection points are E–far from ρ^V_H. By bounded geodesic image, we have, say, $\mathsf{d}_V(a, b) \leqslant E$, a contradiction. A similar argument rules out $H \subsetneq V$. Hence $H \perp V$. □

Lemma 1.26. *Let $A, B \subset \mathcal{X}$ be κ–hierarchically quasiconvex sets. Then there exists $K = K(\kappa, \mathcal{X}, \mathfrak{S})$ so that for all $a \in \mathcal{X}$ we have $\mathsf{d}(a, B) \asymp_{K,K} \mathsf{d}(a, \mathfrak{g}_B(a))$. Moreover, for any $a \in A$:*

$$\mathsf{d}(A, B) \asymp_{K,K} \mathsf{d}(\mathfrak{g}_B(a), \mathfrak{g}_A(\mathfrak{g}_B(a))).$$

Proof. First let $a \in \mathcal{X}$ and $b \in B$. Recall that for $U \in \mathfrak{S}$, the map $p_{U,B} \colon \mathcal{C}U \to \pi_U(B)$ is coarsely the closest-point projection. For any $U \in \mathfrak{S}$, we have $\mathsf{d}_U(a, p_{U,B}(\pi_U(a))) \leqslant \mathsf{d}_U(a, b) + 1$. By the definition of the gate, and the distance formula, we thus have K', depending on κ, so that $\mathsf{d}(a, \mathfrak{g}_B(a)) \leqslant K'\mathsf{d}(a, b) + K'$. Since this holds for any $b \in B$, this proves the first assertion.

Now let $a \in A$ and let $U \in \mathfrak{S}$. Then $p_{U,A}(p_{U,B}(\pi_U(a)))$ lies uniformly close to any $\mathcal{C}U$–geodesic from $\pi_U(a)$ to $p_{U,B}(\pi_U(a))$, so by the distance formula and the definition of the gate, $\mathsf{d}(a, \mathfrak{g}_B(a)) \geqslant \mathsf{d}(\mathfrak{g}_B(a), \mathfrak{g}_A(\mathfrak{g}_B(a)))/K' - K'$ for K' depending only on $\mathcal{X}, \mathfrak{S}$, and κ.

Choose $a \in A$ so that $\mathsf{d}(A, B) \geqslant \mathsf{d}(a, B) - 1$. Then $\mathsf{d}(A, B) \geqslant K'\mathsf{d}(a, g_B(a))/K' - K' - 1$, by the first assertion and the choice of a. As above, $\mathsf{d}(a, \mathfrak{g}_B(a)) \geqslant \mathsf{d}(\mathfrak{g}_B(a), \mathfrak{g}_A(\mathfrak{g}_B(a)))/K' - K'$. Combining these facts shows that, up to uniform constants, $\mathsf{d}(A, B)$ is bounded below by $\mathsf{d}(\mathfrak{g}_B(a), \mathfrak{g}_A(\mathfrak{g}_B(a)))$, as required. □

1.3. **Wallspaces.** *Wallspaces* were introduced by Haglund–Paulin [HP98] and there are now numerous variants of the notion, surveyed in [HW14]. Here, we recall the relevant definitions for Section 2. See, e.g., [HW14] for more background on CAT(0) cube complexes.

Definition 1.27 (Wallspace, coherent orientation)**.** A *wallspace* $(\mathcal{S}, \mathcal{W})$ consists of a set $\mathcal{S}$ and a collection $\mathcal{W} = \{(\overleftarrow{W}, \overrightarrow{W})\}$ of partitions of $\mathcal{S}$; each such partition is called a *wall*. The subsets $\overleftarrow{W}, \overrightarrow{W} \subset \mathcal{S}$ are the *halfspaces associated to* $(\overleftarrow{W}, \overrightarrow{W})$. A *orientation* x of $\mathcal{W}$ is a map $\mathcal{W} \ni (\overleftarrow{W}, \overrightarrow{W}) \mapsto x(\overleftarrow{W}, \overrightarrow{W}) \in \{\overleftarrow{W}, \overrightarrow{W}\}$. The orientation x is *coherent* if $x(\overleftarrow{W}, \overrightarrow{W}) \cap x(\overleftarrow{W}', \overrightarrow{W}') \neq \emptyset$ for all $(\overleftarrow{W}', \overrightarrow{W}'), (\overleftarrow{W}, \overrightarrow{W}) \in \mathcal{W}$. The orientation x is *canonical* if there exists $s \in \mathcal{S}$ so that $s \in x(\overleftarrow{W}', \overrightarrow{W}')$ for all but finitely many $(\overleftarrow{W}', \overrightarrow{W}') \in \mathcal{W}$. When $\mathcal{W}$ is finite, as it is in this paper, any orientation is canonical.

Definition 1.28 (Dual cube complex)**.** The *dual cube complex* $C = C(\mathcal{S}, \mathcal{W})$ associated to the wallspace $(\mathcal{S}, \mathcal{W})$ is the CAT(0) cube complex whose 0–cubes are the coherent, canonical orientations of $\mathcal{W}$, with two 0–cubes joined by a 1–cube if the corresponding orientations differ on exactly one wall. The resulting graph is median [CN05, Nic04, Sag95] and thus the 1–skeleton of a uniquely determined CAT(0) cube complex [Che00] which we call C. Note that, given a CAT(0) cube complex C, each hyperplane W yields a wall in $C^{(0)}$ by partitioning $C^{(0)}$ into the vertex sets of the two components of $C - W$. The cube complex dual to the resulting wallspace is exactly C.

Definition 1.29 (Hyperplane, crossing)**.** A *hyperplane* in C is a connected subspace whose intersection with each cube $c = [-1, 1]^n$ is either $\emptyset$ or a subspace obtained by restricting exactly one coordinate to 0.

The hyperplanes in $C(\mathcal{S}, \mathcal{W})$ correspond bijectively to the walls in $\mathcal{W}$. Moreover, two hyperplanes have nonempty intersection if and only if the corresponding walls *cross* in the sense that all four possible intersections of associated halfspaces are nonempty. It follows that the dimension of C is equal to the largest cardinality of a subset of $\mathcal{W}$ consisting of pairwise-crossing walls.

We occasionally use the *convex hull* of a set $A \subset C(\mathcal{S}, \mathcal{W})$: this is the largest subcomplex contained in the intersection of all halfspaces containing A.

1.4. **Ultralimits and asymptotic cones.** Let (M, d) be a metric space and let $\omega \subset 2^{\mathbb{N}}$ be a non-principal ultrafilter on $\mathbb{N}$. Given a sequence $m = (m_n \in M)_{n\in\mathbb{N}}$ of *observation points* and a positive sequence $s = (s_n)_{n\in\mathbb{N}}$ with $s_n \xrightarrow{n} \infty$, the *asymptotic cone* $\mathbf{M}$ is the ultralimit of the based metric spaces $\lim_\omega(M, m_n, \frac{d}{s_n})$: define a pseudometric $\mathbf{d}$ on $\prod_n M$ by $\mathbf{d}(y, z) = \lim_\omega \frac{d(y_n, z_n)}{s_n}$, and consider the induced pseudometric on the component containing m, i.e.,

$$\widehat{M} = \left\{ (y_n)_{n\in\mathbb{N}} \in \prod_n (M, \frac{d}{d_n}) \ : \ \mathbf{d}(y, m) < \infty \right\}.$$

Then $\mathbf{M}$ is the associated quotient metric space, obtained from $\widehat{M}$ by identifying points y and z for which $\mathbf{d}(y,z) = 0$. We refer the reader to [Dru02] for additional background on asymptotic cones.

We will adopt the following notational conventions. We denote by ω a fixed non-principal ultrafilter on $\mathbb{N}$. Given a sequence $(M_i)_{i\in\mathbb{N}}$ of based metric spaces, we denote by $\mathbf{M}$ the corresponding ultralimit. Given $m \in \mathbf{M}$, a representative of m is a sequence $(m_i \in M_i)_{i\in\mathbb{N}}$, and, when there is no possibility of confusion, we use a boldface letter to denote this representative, viz. $\mathbf{m} = (m_i)$.

We also denote by ${}^{\omega}\mathbb{R}_+$ the ultrapower of the set $\mathbb{R}_+$ of nonnegative reals. Given $\lambda \in {}^{\omega}\mathbb{R}_+$, we sometimes use the notation, e.g., $\mathbf{r}$ to denote a sequence $(r_m)_{m\in\mathbb{N}}$ representing λ.

1.5. **Median, coarse median, quasimedian.** We recall some background on median and coarse median spaces; the reader is referred to [Bow13, Bow15] for a more detailed discussion.

The discussion of coarse median spaces in [Bow13] is given in terms of *(finite) median algebras*. For concreteness, we first consider only the following example of a (finite) median algebra: let $\mathcal{Y}$ be a CAT(0) cube complex (with finitely many 0–cubes). Recall that there exists a *median* map $\mu\colon (\mathcal{Y}^{(0)})^3 \to \mathcal{Y}^{(0)}$ with the property that, for all $x_1, x_2, x_3 \in \mathcal{Y}^{(0)}$, the 0–cube $\mu(x_1,x_2,x_3)$ lies on a combinatorial geodesic from x_i to x_j for all distinct $i,j \in \{1,2,3\}$, see e.g., [Che00]. This 0–cube with the given property is unique.

Remark 1.30 (Median and walls)**.** Let $\mathcal{Y}$ be a CAT(0) cube complex and let x,y,z be 0–cubes. The *median*, $\mu = \mu(x,y,z)$, can be described in terms of orientations of walls as follows. If W is a wall in $\mathcal{Y}$ so that some associated halfspace W^+ contains x,y,z, then μ orients W toward W^+. Otherwise, W has two associated halfspaces $W^{\pm}$ so that W^+ contains exactly two of the points $\{x,y,z\}$ and W^- contains exactly one of these points. Then μ orients W toward W^+. This choice of orientation of all walls is coherent and easily verified to yield a 0–cube which is the median of x,y,z.

The above discussion provides the basis for the definition of a coarse median space.

Definition 1.31 (Coarse median space; [Bow13])**.** Let $(\mathcal{L},\mathsf{d})$ be a metric space and let $\mu\colon \mathcal{L}^3 \to \mathcal{L}$ be a ternary operation. We say that $\mathcal{L}$, equipped with μ, is a *coarse median space* if there exists a constant k and a map $h\colon \mathbb{N} \to [0,\infty)$ so that the following hold:

- For all $x,y,z,x',y',z' \in \mathcal{L}$,
$$\mathsf{d}(\mu(x,y,z),\mu(x',y',z')) \leqslant k(\mathsf{d}(x,x') + \mathsf{d}(y,y') + \mathsf{d}(z,z')) + h(0).$$
- For all $p \in \mathbb{N}$ and $A \subseteq \mathcal{L}$ with $|A| \leqslant p$, there is a CAT(0) cube complex $\mathcal{Y}_A$ with finite 0–skeleton and median map μ_A, and maps $f\colon A \to \mathcal{Y}_A^{(0)}$ and $g\colon \mathcal{Y}_A^{(0)} \to A$ so that the following hold:
  - $\mathsf{d}(\mu(g(x),g(y),g(z)),g(\mu_A(x,y,z))) \leqslant h(p)$ for all $x,y,z \in \mathcal{Y}_A^{(0)}$;
  - $\mathsf{d}(a,g(f(a))) \leqslant h(p)$ for all $a \in A$.

The *coarse median rank* ν of $\mathcal{L}$ is the smallest integer ν so that $\mathcal{Y}_A$ can be taken to have dimension $\leqslant \nu$ for all finite A.

It was shown in [BHS15] that every hierarchically hyperbolic space is a coarse median space; we refer the reader there for details of the construction.

Definition 1.32 (Quasimedian map)**.** Let $\mathcal{Y}$ be a CAT(0) cube complex with median map $\mu_{\mathcal{Y}}$ on its 0–skeleton. Let $(\mathcal{L},\mu,\mathsf{d})$ be a coarse median space. Let $h \geqslant 0$. An *h–quasimedian map* is a map $q\colon \mathcal{Y} \to \mathcal{L}$ for which
$$\mathsf{d}(\mu(q(x),q(y),q(z)),q(\mu_{\mathcal{Y}}(x,y,z))) \leqslant h$$
for all $x,y,z \in \mathcal{Y}$.

Note that quasimedian maps are precisely what [Bow13] calls “quasimorphisms”.

Finally, we recall that a set $\mathcal{M}$ equipped with a ternary operation $\mu\colon \mathcal{M}^3 \to \mathcal{M}$ is a *median algebra* if for all finite $A \subset \mathcal{M}$, there is a finite $B \subset \mathcal{M}$ so that $A \subseteq B$, and B is closed under μ, and (B, μ) is a finite median algebra in the above sense (i.e., we can identify its elements with points in a finite CAT(0) cube complex in such a way that μ coincides with the cubical median). The *rank* of a median algebra is defined as in Definition 1.31 in terms of the dimensions of the cube complexes approximating finite sets.

Given $a, b \in \mathcal{M}$, the *interval* $[a, b]$ is the set of $c \in \mathcal{M}$ with $\mu(a, b, c) = c$, and $\mathcal{N} \subset \mathcal{M}$ is *median convex* if $[a, b] \subseteq \mathcal{N}$ whenever $a, b \in \mathcal{N}$.

If $\mathcal{M}$ is also a Hausdorff topological space, and μ is continuous, then $(\mathcal{M}, \mu)$ is a *topological median algebra*. We consider the following special case. Let (M, d) be a metric space. For any $a, b \in M$, let $[a, b]$ be the set of $c \in M$ for which $\mathsf{d}(a, b) = \mathsf{d}(a, c) + \mathsf{d}(c, b)$. If M has the property that for all $a, b, c \in M$, the intersection $[a, b] \cap [b, c] \cap [c, a]$ consists of a single point $\mu(a, b, c)$, then the map $(a, b, c) \mapsto \mu(a, b, c)$ makes (M, d) a topological median algebra. In this situation, we say M is a *median (metric) space*. The metric notion of an interval agrees with the median notion discussed above.

It is shown in Theorem 2.3 of [Bow13] that any asymptotic cone of a coarse median space of rank ν is a median space of rank ν, where the median of points represented by sequences $(x_n), (y_n), (z_n)$ is represented by a sequence whose n^{th} term is the coarse median of x_n, y_n, z_n.

Definition 1.33 (Block, median gate)**.** Let (M, d) be a median metric space. A *n–block* in M is a median convex subspace isometric to the product of n nontrivial compact intervals, endowed with the ℓ_1 metric.

If $N \subset M$ is a closed median convex subset, a *median gate map* $\mathfrak{g}_N\colon M \to N$ is a map such that $\mathfrak{g}_N(m) \in [m, n]$ for all $m \in M, n \in N$.

If $\mathfrak{g}_N$ exists, then it is unique; if intervals in M are compact, as occurs when M is complete and of finite rank, then $\mathfrak{g}_N$ exists for all closed median convex N. If N, N' are median convex, then $\mathfrak{g}_N(N')$ is again median convex; see [Bow15].

2. Cubulation of hulls

Fix a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$. In this section, we prove that the hull of any finite set $A \subset \mathcal{X}$ can be cubulated. Roughly, our walls are built in the following way. We consider $U \in \mathfrak{S}$ and consider a tree which approximates the convex hull of $\pi_U(A)$ in $\mathcal{C}U$. We then find an appropriate separated net in this tree and, for each point in this net, we use π_U^{-1} of a connected component of the complement as one of our walls.

Specifically, it is the goal of this section to prove:

Theorem 2.1. *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space and let $k \in \mathbb{N}$. Then there exists M_0 so that for all $M \geqslant M_0$ there is a constant C_1 so that for any $A \subset \mathcal{X}$ of cardinality $\leqslant k$, there is a C_1–quasimedian (C_1, C_1)–quasi-isometry $\mathfrak{p}_A\colon \mathcal{Y} \to H_\theta(A)$, where $\mathcal{Y}$ is a CAT(0) cube complex.*

Moreover, let $\mathcal{U}$ be the set of $U \in \mathfrak{S}$ so that $\mathsf{d}_U(x, y) \geqslant M$ for some $x, y \in A$. Then $\dim \mathcal{Y}$ is equal to the maximum cardinality of a set of pairwise-orthogonal elements of $\mathcal{U}$.

Finally, there exist 0–cubes $y_1, \ldots, y_{k'} \in \mathcal{Y}$ so that $k' \leqslant k$ and $\mathcal{Y}$ is equal to the convex hull in $\mathcal{Y}$ of $\{y_1, \ldots, y_{k'}\}$.

The proof is carried out over the next several subsections. We fix once and for all $(\mathcal{X}, \mathfrak{S})$, some $k \in \mathbb{N}$, and a subset $A = \{x_1, \ldots, x_k\} \subseteq \mathcal{X}$.

2.1. The candidate finite CAT(0) cube complex.

Fix $U \in \mathfrak{S}$. For each $x_j \in A$, recall that $\pi_U(x_j)$ is a subset of the δ–hyperbolic space $\mathcal{C}U$ of diameter at most E; for each j, choose $\ell_j^U \in \pi_U(x_j)$, to obtain k points $\ell_1^U, \ldots, \ell_k^U \in \mathcal{C}U$. There exists $C = C(k, \delta)$ so that there is a finite tree T_U and an embedding $T_U \hookrightarrow \mathcal{C}U$, sending edges to geodesics of $\mathcal{C}U$, with the following properties:

- $\mathsf{d}_U(p, q) \leqslant \mathsf{d}_{T_U}(p, q) \leqslant \mathsf{d}_U(p, q) + C$ for all $p, q \in T_U$;
- ℓ_j^U is a leaf of T_U for $1 \leqslant j \leqslant k$;
- each leaf of T_U lies in $\{\ell_1^U, \ldots, \ell_k^U\}$.

This is the usual spanning tree of a finite subset of a hyperbolic space; see [Gro87]. The given properties of T_U ensure that, up to increasing C uniformly, $\mathsf{d}_{haus}(T_U, \mathrm{hull}_{\mathcal{C}U}(\pi_U(A))) \leqslant C$.

Our choice of T_U ensures that, for each $x_j \in A$, $\pi_U(x_j) \subset \mathcal{C}U$ contains a leaf of T_U, and every leaf of T_U is contained in $\pi_U(x_j)$ for some $x_j \in A$.

Let M be a (large) constant to be specified below. We will point out the conditions that M must satisfy as we proceed. Let $\mathcal{U}$ be the (finite) set of all $U \in \mathfrak{S}$ with $\mathsf{diam}(\pi_U(A)) \geqslant 100M$. Let $\mathcal{U}_1 \subseteq \mathcal{U}$ be the set of $\sqsubseteq$–minimal elements of $\mathcal{U}$. Given $\mathcal{U}_{n-1}$, let $\mathcal{U}_n \subseteq \mathcal{U}$ be the set of all $\sqsubseteq$–minimal elements of $\mathcal{U} - \mathcal{U}_{n-1}$. Finite complexity ensures that there is some s so that $\bigcup_{n=1}^s \mathcal{U}_s = \mathcal{U}$. For each $U \in \mathcal{U}$, let $\mathcal{U}^{\sqsubseteq, U} = \{V \in \mathcal{U} : V \subsetneq U\}$. For each $V \in \mathcal{U}^{\sqsubseteq, U}$, choose $r_U^V \in T_U$ closest to ρ_U^V; the set of choices is bounded diameter (moreover, in Lemma 2.4, we prove that r_U^V is $100EC$–close to ρ_U^V).

Starting with each $U \in \mathcal{U}_1$ and then repeating for $\mathcal{U}_2$ up to $\mathcal{U}_s$, we choose a finite set of elements $p_i^U \in T_U$ satisfying the following conditions (which provide that the p_i^U together with the r_U^V provide a $10M$–net which is M–separated):

1. $\mathsf{d}_U(p_i^U, x_j) \geqslant M$,
2. $\mathsf{d}_U(p_i^U, p_j^U) \geqslant M$,
3. $\mathsf{d}_U(p_i^U, r_U^V) \geqslant M$ for each $V \in \mathcal{U}^{\sqsubseteq, U}$ (when $U \in \mathcal{U}_1$, there are no such V), and
4. each component of $T_U - \Big(\{p_i^U\} \cup \{r_U^V\}_{V \in \mathcal{U}^{\sqsubseteq, U}}\Big)$ has diameter at most $10M$ (when $U \in \mathcal{U}_1$, there are no such V, so the criterion is only about complements of the $\{p_i^U\}$).

For each $U \in \mathfrak{S}$, let β_U be the composition of π_U and a closest point projection to T_U (for each $p \in H_\theta(A)$, we have $\mathsf{diam}_{\mathcal{C}U}(\pi_U(p) \cup \beta_U(p)) \leqslant 10(E + \theta + C)$).

Definition 2.2 (Walls in $H_\theta(A)$)**.** Given $U \in \mathcal{U}$ and $\{p_i^U\}$ as above, for each i we define a partition $H_\theta(A) = \overleftarrow{W}_i^U \sqcup \overrightarrow{W}_i^U$ of $H_\theta(A)$ as follows. Choose a component T_U' of $T_U - \{p_i^U\}$ and let $\overleftarrow{W}_i^U = \beta_U^{-1}(T_U') \cap H_\theta(A)$, and set $\overrightarrow{W}_i^U = H_\theta(A) - (\overleftarrow{W}_i^U)$. Let $\mathcal{L}_i^U = (\overleftarrow{W}_i^U, \overrightarrow{W}_i^U)$.

Observe that the (finite) set of walls in $H_\theta(A)$ specified in Definition 2.2 depends on our choice of M (since that determines $\mathcal{U}$) and on our choice of the p_i^U (which is also constrained by the choice of M). Let $\mathcal{Y}$ be the CAT(0) cube complex dual to the wallspace just defined. Since the set of walls is finite, there is exactly one 0–cube in $\mathcal{Y}$ for each coherent orientation of all the walls (recall that a coherent orientation is a choice of halfspace for each wall such that, for any two walls, the chosen halfspaces have nonempty intersection).

2.2. Lemmas supporting consistency of certain tuples.

Lemma 2.3. *For all $M > 10E$, the following holds. Let $U \in \mathcal{U}$ and $V \in \mathfrak{S}$. If $U \pitchfork V$ then ρ_V^U is E–close to some $\pi_V(x_i)$, and hence $2E$–close to T_V.*

Proof. Since $U \in \mathcal{U}$, we have $\mathsf{diam}_{\mathcal{C}U}(\pi_U(A)) \geqslant 100M > 10^3E$. Hence we can choose $x_i \in A$ so that $\mathsf{d}_U(x_i, \rho_U^V) > E$. Consistency yields $\mathsf{d}_V(x_i, \rho_V^U) \leqslant E$. Since $\pi_V(x_i)$ has diameter $\leqslant E$ and contains a leaf of T_V, we have $\mathsf{d}_V(T_V, \rho_V^U) \leqslant 2E$. □

Lemma 2.4. *For any $M > 10E$, the following holds. Let $U \in \mathcal{U}, V \in \mathfrak{S}$, with $U \subsetneq V$. Then $\mathsf{d}_V(\rho_V^U, T_V) \leqslant 100EC$.*

Proof. Suppose that $\mathsf{d}_V(\rho^U_V, T_V) > 100EC$. Then, since T_V C–coarsely coincides with $\mathrm{hull}_{\mathcal{C}V}(A)$, and the latter is $5E$–quasiconvex, we have that ρ^U_V lies at distance greater than E from any geodesic joining points in $\pi_V(A)$. Hence, by consistency and bounded geodesic image, any such geodesic projects to a geodesic in $\mathcal{C}U$ of diameter at most E, i.e., $\pi_U(A)$ has diameter bounded by $10E$. This contradicts $U \in \mathcal{U}$, provided $M > 10E$. □

Lemma 2.5. *For any $M > 10E$ the following holds. Consider $U \in \mathcal{U}$ and any $V \in \mathfrak{S}$ with $V \subsetneq U$. Then for each $x \in T_U - \mathcal{N}_M(\rho^V_U)$ there exists $x_j \in A$ with $\mathsf{d}_V(\rho^U_V(x), x_j) \leqslant 2E$ (in particular, $\rho^U_V(x)$ is $10E$–close to T_V).*

Proof. There exists a leaf of T_U, contained in $\pi_U(x_j)$ for some $x_j \in A$, in the same connected component of $T_U - \mathcal{N}_{M/2}(\rho^V_U)$ as x. Geodesics from x to $\pi_U(x_j)$ thus stay E–far from ρ^V_U, so that the desired conclusion follows from bounded geodesic image (and consistency, which says $\mathsf{diam}_V(\pi_V(x_j) \cup \rho^U_V(\pi_U(x_j))) \leqslant E$). □

2.3. **The proof of Theorem 2.1.** We now prove Theorem 2.1. Some auxiliary lemmas appear immediately below the proof, organized according to which part of the proof they support.

Proof of Theorem 2.1. We break the proof into several parts.

Definition of $\mathfrak{p}_A$: We first define $\mathfrak{p}_A \colon \mathcal{Y} \to \mathcal{X}$, noting that it suffices to define $\mathfrak{p}_A$ on the 0–skeleton of $\mathcal{Y}$. Let $p \in \mathcal{Y}^{(0)}$; we view p as a coherent orientation of the walls $\mathcal{L}^U_i$ provided by Definition 2.2.

For $U \in \mathcal{U}$, $V \in \mathfrak{S}$ and each p^U_i (which we recall gives a pair $\{\overleftarrow{W}^U_i, \overrightarrow{W}^U_i\}$), we can consider $\overline{W}_i(U) \in \{\overleftarrow{W}^U_i, \overrightarrow{W}^U_i\}$ which is the halfspace given by the orientation p, namely $p(\overleftarrow{W}_i(U), \overrightarrow{W}_i(U))$. We let $S_{U,i,V}(p) \subseteq T_V$ be the convex hull in T_V of $\beta_V(\overline{W}_i(U))$, where, as above, β_V is the composition of projection to $\mathcal{C}V$ and the closest point projection to T_V.

By the definition of a coherent orientation, for any U, i, U', i', we have $\beta_V(\overline{W}_i(U)) \cap \beta_V(\overline{W}_{i'}(U')) \neq \varnothing$, whence $S_{U,i,V}(p) \cap S_{U',i',V}(p) \neq \varnothing$. The Helly property for trees thus ensures that $\bigcap_{U,i} S_{U,i,V}(p) \neq \varnothing$ for each $V \in \mathfrak{S}$, and we let $b_V = b_V(p) = \bigcap_{U,i} S_{U,i,V}(p)$. Lemma 2.8, below, proves that $\mathsf{diam}(b_V)$ are uniformly bounded. Lemma 2.9, below, shows the (b_V) are η–consistent, where $\eta = \eta(M, k, \mathcal{X})$.

We can now define $\mathfrak{p}_A(p) \in \mathcal{X}$ to be a realization point associated to (b_U) via Theorem 1.4. Specifically, there exists $\xi = \xi(\eta, E)$ so that for all $U \in \mathfrak{S}$, we have $\mathsf{d}_U(\pi_U(\mathfrak{p}_A(p)), b_U) \leqslant \xi$.

The image of $\mathfrak{p}_A$ coarsely coincides with $H_\theta(A)$: For any $x \in H_\theta(A)$, one can orient the walls coherently by choosing, for each wall, the halfspace containing x. The resulting 0–cube $p \in \mathcal{Y}$ has the property that $\mathsf{d}_{\mathcal{X}}(x, \mathfrak{p}_A(p)) \leqslant C'_1$, where $C'_1 = C'_1(M, k, \mathcal{X})$. Hence $H_\theta(A)$ lies in a uniform neighborhood of $\mathrm{im}\, \mathfrak{p}_A$. On the other hand, if $p \in \mathcal{Y}$, then $\pi_U(\mathfrak{p}_A(p))$ lies uniformly close to $\mathrm{hull}(\pi_U(A))$, so hierarchical quasiconvexity of $H_\theta(A)$ ensures that $\mathfrak{p}_A(p)$ lies uniformly close to $H_\theta(A)$, i.e., $\mathrm{im}\, \mathfrak{p}_A$ lies in a uniform neighborhood of $H_\theta(A)$.

Distance estimates: For $p \in \mathcal{Y}$, we say p^U_i is a *separator* for p if p^U_i separates $\beta_U(\mathfrak{p}_A(p))$ from b_U. We call U the *support* of the separator. In Lemma 2.11 we prove there is a uniform bound, T, so that for each $p \in \mathcal{Y}$ there are at most T separators for p.

We now relate the number of walls separating a pair of points in $\mathcal{Y}$ to the number of points separating their images under $\mathfrak{p}_A$. Namely, if $p, q \in \mathcal{Y}$, then $\mathsf{d}_{\mathcal{Y}}(p, q)$ is the number of walls between p and q, which in turn is the sum of the numbers of p^V_i separating $b_V(p)$ from $b_V(q)$, as V varies. By Lemma 2.11, up to an additive error this is the same as the sum over V of the number of p^V_i separating $\beta_V(\mathfrak{p}_A(p))$, $\beta_V(\mathfrak{p}_A(q))$; we write $Q(p, q)$ to denote this sum.

Observe that: if, for some V, there exist distinct $p^V_i, p^V_{i'}$ separating $\beta_V(\mathfrak{p}_A(p))$ from $\beta_V(\mathfrak{p}_A(q))$, then V contributes to the distance formula sum between p and q, at some fixed threshold L chosen in terms of E. Moreover, V also contributes to the distance formula sum in the case where $\beta_V(\mathfrak{p}_A(p)), \beta_V(\mathfrak{p}_A(q))$ are both C–close to $\pi_V(A)$ and there exists at least one p_i^V separating $\beta_V(\mathfrak{p}_A(p)), \beta_V(\mathfrak{p}_A(q))$.

Applying Lemma 2.6 and Lemma 2.10, we have

$$\mathsf{d}_{\mathcal{X}}(\mathfrak{p}_A(p), \mathfrak{p}_A(q)) \asymp \sum_{U \in \mathfrak{S}} \{\!\{\mathsf{d}_U(\mathfrak{p}_A(p), \mathfrak{p}_A(q))\}\!\}_L \geqslant Q(p,q) - 100EC\theta N,$$

where N is the constant from Lemma 2.10. Hence there exists $C_1'' = C_1''(M, \mathcal{X}, k)$ so that $\mathsf{d}_{\mathcal{X}}(\mathfrak{p}_A(p), \mathfrak{p}_A(q)) \geqslant \mathsf{d}_{\mathcal{Y}}(p,q)/C_1'' - C_1''$ for $p, q \in \mathcal{Y}$.

$\mathfrak{p}_A$ is coarsely Lipschitz: Crossing one hyperplane of $\mathcal{Y}$ corresponds to changing only one coordinate (b_U) as above by a bounded amount, so there exists $C_1''' = C_1'''(M, k, \mathcal{X})$ so that $\mathfrak{p}_A$ is (C_1''', C_1''')–coarsely Lipschitz.

Dimension: The assertion about dimension follows from Lemma 2.13 and the well-known fact that any finite set of n pairwise crossing hyperplanes in a CAT(0) cube complex intersect in the barycenter of some n–cube.

Convex hull: For each $x_j \in A$, let y_j be the orientation of the walls in $H_\theta(A)$ obtained by choosing, for each wall $(\overleftarrow{W}_i^U, \overrightarrow{W}_i^U)$, the halfspace containing x_j. This orientation is coherent by definition, so determines a 0–cube of $\mathcal{Y}$, which we also denote y_j. By construction, each wall separates two elements of A, so every hyperplane of $\mathcal{Y}$ separates two of the chosen 0–cubes y_i, y_j. Thus no intersection of combinatorial halfspaces properly contained in $\mathcal{Y}$ contains all of the y_j, so $\mathcal{Y}$ is the convex hull in $\mathcal{Y}$ of the set of y_j.

Conclusion: Lemma 2.7 provides C_1'''' so that $\mathfrak{p}_A$ is C_1''''–quasimedian, so the proof is complete once we take $C_1 = \max\{C_1', C_1'', C_1''', C_1''''\}$. ☐

Lemma 2.6. *Let $U \in \mathcal{U}$. For each $x, y \in H_\theta(A)$, we have $\mathsf{d}_U(x,y) + 50EC\theta \geqslant |\{i : p_i^U \in [\beta_U(x), \beta_U(y)]\}|$. Moreover, if $\pi_U(x), \pi_U(y)$ are both C–close to $\pi_U(A)$, then $\mathsf{d}_U(x,y) \geqslant |\{i : p_i^U \in [\beta_U(x), \beta_U(y)]\}|$.*

Proof. Let $x, y \in H_\theta(A)$. Recall that $\mathsf{diam}(\pi_U(x) \cup \beta_U(x)) \leqslant 10(E + C + \theta)$, so $\mathsf{d}_U(x,y) \geqslant \mathsf{d}_U(\beta_U(x), \beta_U(y)) - 20(E + C + \theta)$. Hence $\mathsf{d}_U(x,y) \geqslant \mathsf{d}_{T_U}(\beta_U(x), \beta_U(y)) - 40EC\theta$. Hence $\mathsf{d}_U(x,y) \geqslant |\{i : p_i^U \in [\beta_U(x), \beta_U(y)]\}| - 40EC\theta - 1$, as required. The "moreover" statement follows in a similar way using the fact that the p_U^i are M–far from leaves of T_U. ☐

Lemma 2.7. *There exists $C_1'''' = C_1''''(\mathcal{X}, k, M)$ so that $\mathfrak{p}_A$ is C_1''''–quasimedian.*

Proof. Let $\mu \colon \mathcal{X}^3 \to \mathcal{X}$ be the coarse median map.

Let $x, y, z \in \mathcal{Y}$, and let m be their median. By Remark 1.30, m corresponds to the following orientation of the walls of $\mathcal{Y}$: for each wall W, $m(W)$ is the halfspace which contains at least two of x, y, z. In other words, for each $U \in \mathcal{U}$ and each $p_i^U \in T_U$, the orientation m assigns to $\{\overleftarrow{W}_i(U), \overrightarrow{W}_i(U)\}$ is the halfspace $\overline{W}_i(U)$ assigned by at least two of the orientations x, y, z.

By definition, for any $V \in \mathfrak{S}$, we have $b_V(m) = \bigcap_{U \in \mathcal{U}, i} S_{U,i,V}(m)$, where, for each U, i, we have that $S_{U,i,V}(m)$ coincides with at least two of $S_{U,i,V}(x), S_{U,i,V}(y), S_{U,i,V}(z)$.

In particular, for each $V \notin \mathcal{U}$, we have that $b_V(m)$ coarsely coincides with each of $\beta_V(x), \beta_V(y), \beta_V(z)$.

Also, for each $U \in \mathcal{U}$ and each p_i^U, we have that $b_U(m)$ lies in the same p_i^U–halfspace of T_U as at least two of the points $b_U(x), b_U(y), b_U(z)$. Hence $b_U(m)$ lies in the same p_i^U–halfspace of T_U as m_U, where m_U is the median of $b_U(x), b_U(y), b_U(z)$ in the tree T_U. We have shown that no p_i^U separates $b_U(m)$ from m_U, for any $U \in \mathcal{U}$.

Our $(1, C)$–quasi-isometrically embedded choice of T_U ensures that m_U is, up to uniformly bounded error, a coarse median point for the images in $\mathcal{C}U$ of $\mathfrak{p}_A(x), \mathfrak{p}_A(y), \mathfrak{p}_A(z)$. In other words, $\mu(\mathfrak{p}_A(x), \mathfrak{p}_A(y), \mathfrak{p}_A(z))$ is a realization point for $(m_V)_{V \in \mathfrak{S}}$. As shown earlier in the proof of Theorem 2.1, the image of $\mathfrak{p}_A$ coarsely coincides with $H_\theta(A)$, which is hierarchically quasiconvex by Proposition 1.17. Hence $\mu(\mathfrak{p}_A(x), \mathfrak{p}_A(y), \mathfrak{p}_A(z))$ uniformly coarsely coincides with $\mathfrak{p}_A(q)$ for some $q \in \mathcal{Y}$.

The distance estimate in the proof of Theorem 2.1 shows that

$$\mathsf{d}_{\mathcal{X}}(\mathfrak{p}_A(m), \mu(\mathfrak{p}_A(x), \mathfrak{p}_A(y), \mathfrak{p}_A(z))) \asymp \mathsf{d}_{\mathcal{X}}(\mathfrak{p}_A(m), \mathfrak{p}_A(q)) \asymp \mathsf{d}_{\mathcal{Y}}(m, q)$$

can be bounded in terms of the number of walls separating m, q. Up to additive error, this is just the sum over $U \in \mathcal{U}$ of the number of p_i^U separating $b_U(m)$ from m_U, which we established above was 0, as required. □

2.3.1. *Lemmas supporting realization.* The first two lemmas are used to construct a point in $\mathcal{X}$ via realization.

Lemma 2.8. *There exists $\tau = \tau(M, k)$ (independent of V) so that $\mathsf{diam}(b_V) \leqslant \tau$.*

Proof. If $V \in \mathfrak{S} - \mathcal{U}$, then $\mathsf{diam}(b_V) \leqslant \mathsf{diam}(T_V) \leqslant 100M$. Hence suppose that $V \in \mathcal{U}$.

By definition of the p_i^V, there exists $\tau = \tau(M, k) \geqslant 50M$ so that for all $\beta_V(x), \beta_V(y) \in T_V$ satisfying $\mathsf{d}_V(x, y) > \tau$, there exists $\alpha \in \{p_i^V\}_i \cup \{r_V^W\}_{W \in \mathcal{U}_1 \cap V^{\sqsubseteq}, \mathcal{U}}$ so that α is $10M$–far from x, y and from all points of T_V of valence larger than 2. The restriction to $\mathcal{U}_1$ is justified by the fact that for $W' \subsetneq W \subsetneq U$, we have that $\rho_V^{W'}$ coarsely coincides with ρ_V^W.

Choose any $x, y \in \mathcal{X}$ projecting M–close to b_V, and suppose by contradiction that $\mathsf{d}_V(\beta_V(x), \beta_V(y)) > \tau$. Let α be as above.

If $\alpha = p_i^V$, then we clearly have a contradiction since b_V is contained in one of the connected components of $T_V - \{p_i^V\}$. If $\alpha = r_V^W$, then we write $A \cup \{x, y\} = A' \sqcup A''$, where we group together all elements of $A \cup \{x, y\}$ corresponding to a point of T_V in a given connected component of $T_V - \{r_V^W\}$. By bounded geodesic image and the fact that r_V^W is close to ρ_V^W (Lemma 2.4), $\pi_W(A')$ and $\pi_W(A'')$ are uniformly bounded, so that T_W consists of two uniformly bounded sets, respectively containing $\pi_W(A')$ and $\pi_W(A'')$, that are joined by a segment in T_W which is a geodesic γ of $\mathcal{C}W$ containing no valence–> 2 vertex. Moreover, this geodesic has $\beta_W(x), \beta_W(y)$ uniformly close to its endpoints.

Since $W \in \mathcal{U}_1$, there exists some p_i^W in T_W. Let us show that $S_{W,i,V}(p)$ is far from one of $\beta_V(x)$ or $\beta_V(y)$, which is a contradiction. If there is a p_i^W in T_W, then since p_i^W was chosen far from the leaves of T_W, we have that $p_i^W \in \gamma$, lying at distance $M/2$ from $\beta_W(x)$ from $\beta_W(y)$.

Let $\overline{T}$ be one of the two connected components of $T_W - \{p_i^W\}$. Then $\beta_W^{-1}(\overline{T})$ cannot contain points x', y' with $\beta_V(x'), \beta_V(y')$ far from r_V^W and in different components of $T_V - \{r_V^W\}$, which is the required property of $S_{W,i,V}(p)$. Indeed, otherwise bounded geodesic image would imply that x', y' project respectively close to $\pi_W(A')$ and $\pi_W(A'')$, thus on opposite sides of p_i^W. □

Lemma 2.9. *(b_V) is η–consistent, where $\eta = \eta(M, k, \mathcal{X})$.*

Proof. Let $U \pitchfork V$. If $U, V \in \mathfrak{S} - \mathcal{U}$, we are done because the corresponding coordinates b_U, b_V $(100M + E)$–coarsely coincide with those of, say, x_1. If $U \in \mathcal{U}$ and $V \in \mathfrak{S} - \mathcal{U}$, then any point in $H_\theta(A)$ projects in $\mathcal{C}V$ E–close to T_V and hence $10E$–close to ρ_V^U by Lemma 2.3, so we are done. Now suppose that $U, V \in \mathcal{U}$. Let c_U be a point in T_U $10E$–close to ρ_U^V, and define c_V similarly (c_U and c_V are provided by Lemma 2.3). If both b_U and b_V are $100M$–far from the corresponding ρ, then there are $S_{W,i,U}(p), S_{W',i',V}(p)$ containing b_U, b_V but far from c_U, c_V. There cannot be $q \in \mathcal{X}$ with $\beta_U(q) \in S_{W,i,U}(p), \beta_V(q) \in S_{W',i',V}(p)$ by consistency, implying that the intersection of the halfspaces chosen from $\mathcal{L}_i^W, \mathcal{L}_{i'}^{W'}$ is empty. This contradicts the coherence of the orientation defining p.

Let $U \subsetneq V$. If $V \in \mathfrak{S} - \mathcal{U}$, then by Lemma 2.4 we have that ρ_V^U is $100EC$–close to any point in T_V, in particular b_V. Hence, we can assume $V \in \mathcal{U}$. If $U \in \mathfrak{S} - \mathcal{U}$, similarly,

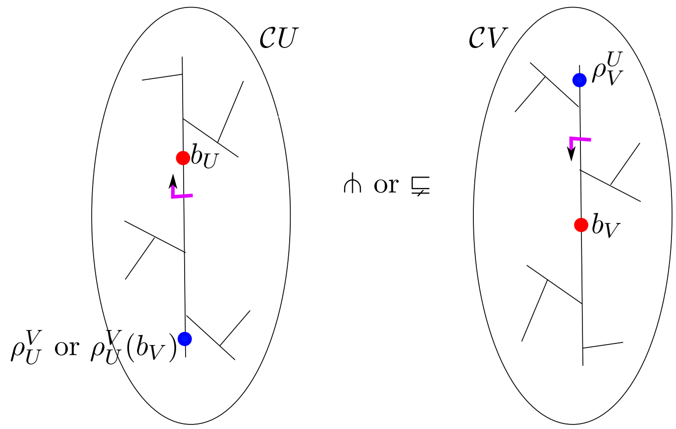

FIGURE 2. Proof of Lemma 2.9. $S_{W,i,U}(p), S_{W',i',V}(p)$ are shown as oriented halfspaces in the trees T_U, T_V.

the corresponding coordinates b_U, b_V coarsely coincide with those of a point in $H_\theta(A)$ that projects close to b_V in $\mathcal{C}V$.

Finally, suppose $U, V \in \mathcal{U}$. The argument is very similar to the final argument in the transverse case above. Let $c_V = r_V^U$ (which is $10E$–close to ρ_V^U by Lemma 2.4); and, as given by Lemma 2.5, we let c_U be a point in T_U which is $100EC$–close to $\rho_U^V(b_V)$. If both b_U and b_V are $100M$–far from the corresponding ρ, then there exist $S_{W,i,U}(p), S_{W',i',V}(p)$ containing b_U, b_V but far from c_U, c_V. By bounded geodesic image, $\rho_U^V(S_{W',i',V}(p))$ has uniformly bounded diameter. Hence, there cannot be $q \in \mathcal{X}$ with $\beta_U(q) \in S_{W,i,U}(p), \beta_V(q) \in S_{W',i',V}(p)$ by consistency, implying that the intersection of the halfspaces chosen from $\mathcal{L}_i^W, \mathcal{L}_{i'}^{W'}$ is empty. This contradicts the coherence of the orientation defining p. □

2.3.2. *Control of separators.* The next two lemmas prove that there is a uniform bound on the number of separators for each point; the first is a straightforward application of Ramsey theory.

Lemma 2.10. *There exists $N = N(\mathcal{X})$ so that for each $x \in H_\theta(A)$ there are at most N elements $U \in \mathcal{U}$ so that $\mathsf{d}_U(\beta_U(x), \pi_U(A)) > 100E$.*

Proof. One axiom of an HHS is that there is a bound, c, on the cardinality of subsets of $\mathfrak{S}$ whose elements are pairwise $\sqsubseteq$–comparable. By [BHS15, Lemma 2.1], c also bounds the maximum cardinality of a set of pairwise orthogonal elements. Given x, consider the set of $U \in \mathfrak{S}$ such that $\mathsf{d}_U(x, A) > 100E$. Ramsey's theorem provides N (the Ramsey number $R(c,c)$) for which either there are at most N such U, or there exist U_1, U_2 with $U_1 \pitchfork U_2$ and $\mathsf{d}_{U_l}(x, A) > 100E$ for $l = 1, 2$. By Lemma 2.3, $\rho_{U_2}^{U_1}$ is $10E$–close to an element of $\pi_{U_2}(A)$ and thus $90E$–far from $\pi_{U_2}(x)$. The same holds with U_1 and U_2 reversed, contradicting consistency. □

Lemma 2.11. *There exists T such that for any $p \in \mathcal{Y}$ there exist at most T separators for p.*

Proof. For any $V \in \mathfrak{S}$, since $\mathsf{d}_V(\mathfrak{p}_A(p), b_V) \leqslant \xi$, the number of separators with support V is bounded in terms of ξ. Hence, by Lemma 2.10, for any $N' \in \mathbb{N}$, there exists $T(N')$ so that, if there are more than $T(N')$ separators for p, then there are at least N' distinct $U \in \mathcal{U}$ so that, for some j, k:

- $\beta_U(\mathfrak{p}_A(p))$ is $100E$–close to $\pi_U(x_j)$;
- there exists a separator p_i^U for p, with support U, separating $\beta_U(x_k)$ from $\beta_U(x_j)$.

The domains U as above are (j,k)–*separators*. Note that if U is a (j,k)–separator, then $\mathsf{d}_U(x_j,x_k) > M > 10E$. Hence, if N' exceeds the constant $N_0 = N_0(100M)$ provided by Lemma 1.6, and the number of separators exceeds $T(N')$, then Lemma 1.6 provides (j,k)–separators U_1, U_2, both properly nested into some V for which $\mathsf{d}_V(r_V^{U_1}, r_V^{U_2}) > 10E + \xi$.

For $l = 1,2$, there exists $p_{i(l)}^{U_l}$ separating a point close to $\pi_{U_l}(x_j)$ from $\pi_{U_l}(x_k)$, so $\mathsf{d}_{U_l}(x_j,x_k) > M$. Hence bounded geodesic image implies that the geodesic in T_V from $\beta_V(x_j)$ to $\beta_V(x_k)$ must pass through $r_V^{U_1}$ and $r_V^{U_2}$. Bounded geodesic image and consistency imply that $\beta_V(\mathfrak{p}_A(p))$ lies E–close to the connected component of $T_V - \{r_V^{U_l}\}$ containing $\pi_V(x_j)$, and the same holds for b_V and $\pi_V(x_k)$. Thus $\mathsf{d}_V(\beta_V(\mathfrak{p}_A(p)), b_V) > \mathsf{d}_V(r_V^{U_1}, r_V^{U_2}) - 2E > \xi$, contradicting the definition of $\mathfrak{p}_A(p)$. □

2.4. **Walls cross if and only if orthogonal.**

Lemma 2.12. *Suppose $U, V \in \mathcal{U}$ and $U \perp V$, and fix any $p \in \mathrm{hull}_{\mathcal{C}U}(A)$, $q \in \mathrm{hull}_{\mathcal{C}V}(A)$. Then there exists $x \in H_\theta(A)$ that coarsely projects to p in $\mathcal{C}U$ and to q in $\mathcal{C}V$.*

Proof. By partial realization, there exists $x' \in \mathcal{X}$ projecting E–close to p in $\mathcal{C}U$ and q in $\mathcal{C}V$. Up to replacing E with a uniform constant depending on θ, the projection $\mathfrak{g}_{H_\theta(A)}(x')$ to $H_\theta(A)$ has the same property, as required. □

Lemma 2.13 (Cross iff orthogonal)**.** *The walls $\mathcal{L}_i^U$ and $\mathcal{L}_j^V$ cross if and only if $U \perp V$.*

Proof. If $U \perp V$, then $\mathcal{L}_i^U$ crosses $\mathcal{L}_j^V$ (recall that this means that each of the four possible intersections of halfspaces, one associated to each wall, is nonempty) by Lemma 2.12.

Conversely, suppose $U \not\perp V$. We claim $\mathcal{L}_i^U$ and $\mathcal{L}_j^V$ do not cross. First, suppose $U \pitchfork V$. Then, by Lemma 2.3, ρ_U^V and ρ_V^U are uniformly close to leaves in the corresponding trees and hence far from p_j^V, p_i^U. Thus, we can choose a halfspace from $\mathcal{L}_i^U$ (resp. $\mathcal{L}_j^V$) so that all its points project far from ρ_U^V (resp. ρ_V^U). The chosen halfspaces are disjoint by consistency. Second, if $U \subsetneq V$, apply the same argument, except that now p_j^V is far from ρ_V^U by construction. □

2.5. **Application to coarse median rank and hyperbolicity.**

In [BHS15, Theorem 7.3], we showed that any HHS is a coarse median space (in the sense of [Bow13]) of rank bounded by the complexity. In the asymphoric case, the following strengthens that result.

Corollary 2.14. *Suppose that $\mathcal{X}$ is asymphoric. Then any cube complex $\mathcal{Y}$ from Theorem 2.1 satisfies* $\dim \mathcal{Y} \leqslant \nu$*, where ν is the rank of $\mathcal{X}$.*

Corollary 2.15. *If $\mathcal{X}$ is an asymphoric HHS of rank ν, then $\mathcal{X}$ is coarse median of rank ν.*

Proof of Corollary 2.14 and Corollary 2.15. Choose M as in the proof of Theorem 2.1; since $M > E$, in particular M exceeds the asymphoricity constant. For any finite $A \subset \mathcal{X}$, let $\mathcal{Y}$ be the cube complex and $\mathcal{Y} \to H_\theta(A)$ be the C_1–quasimedian (C_1, C_1)–quasi-isometry provided by Theorem 2.1. By Lemma 2.13, $\dim \mathcal{Y}$ is equal to the maximal cardinality of sets of pairwise-orthogonal elements of $\mathcal{U}$. But since elements of $\mathcal{U}$ have associated hyperbolic spaces of diameter $\geqslant M$, such subsets have cardinality bounded by ν. This proves Corollary 2.14. Moreover, $\mathcal{Y}^{(0)} \to H_\theta(A)$ is a quasimedian map from a finite median algebra satisfying the condition $(C2)$ from the definition of a coarse median space in [Bow13, Section 8]. The rank of this median algebra is, by definition, $\dim \mathcal{Y} \leqslant \nu$. Hence $\mathcal{X}$ is coarse median of rank ν. □

We can also use the proof of Corollary 2.15 to characterize hyperbolic HHS. We say that a quasi-geodesic metric space X is *hyperbolic* if there exists D so that

- any pair of points of X is joined by a (D,D)–quasi-geodesic, and
- (D,D)–quasi-geodesic triangles are D–thin.

For us, the distinction between hyperbolic geodesic spaces and hyperbolic quasi-geodesic spaces does not matter. Indeed, any quasi-geodesic metric space X is quasi-isometric to a geodesic metric space Y (in fact, a graph). If, in addition, X is hyperbolic then Y is hyperbolic (in the usual sense). There is a number of ways to see this, one of which is the "guessing geodesics" criterion for hyperbolicity from [MS13, Section 3.13][Bow14, Proposition 3.1]. It thus follows from [Bow13, Theorem 2.1] that a coarse median quasigeodesic space is hyperbolic if and only if it has rank at most 1.

We thus get a characterization of HHS which are hyperbolic:

Corollary 2.16. *Let $(\mathcal{X},\mathfrak{S})$ be an HHS. Then the following are equivalent:*

- *$\mathcal{X}$ is coarse median of rank $\leqslant 1$, and is thus hyperbolic;*
- *(Bounded orthogonality) There exists $q \in \mathbb{R}$ so that $\min\{\mathsf{diam}(\mathcal{C}U), \mathsf{diam}(\mathcal{C}V)\} \leqslant q$ for all $U, V \in \mathfrak{S}$ satisfying $U \perp V$.*

Proof. The fact that hyperbolicity implies bounded orthogonality easily follows from the construction of standard product regions. The reverse implication follows from Corollary 2.15, with $\nu = 1$, and the aforementioned [Bow13, Theorem 2.1]. □

Remark 2.17. One can prove that bounded orthogonality implies hyperbolicity using the guessing geodesics criterion instead of the coarse median rank. More specifically, triangles of hierarchy paths are thin because any such triangle is contained in the hull of the vertices, which is quasi-isometric to a 1–dimensional cube complex, i.e. a tree.

3. Quasiflats and asymptotic cones

3.1. **Ultralimits of hulls.** For any hierarchically quasiconvex $A \subseteq \mathcal{X}$ and any $p, q \in A$, $x \in \mathcal{X}$, the coarse median of (p, q, x) lies uniformly close to A. This easily yields:

Lemma 3.1. *For any κ, the ultralimit of any sequence of κ–hierarchically quasiconvex subspaces is median convex.*

Given m, m' in a median space M, we let $\mathrm{hull}(m, m')$ denote the set of $z \in M$ for which the median of m, m', z is z. (Note that $\mathrm{hull}(m, m') = [m.m']$.)

Lemma 3.2. *Let $\boldsymbol{x}, \boldsymbol{y} \in \boldsymbol{\mathcal{X}}$. Then* $\mathrm{hull}(\{\boldsymbol{x}, \boldsymbol{y}\}) = \lim_\omega H_\theta(\{x_n, y_n\})$.

Proof. If $z_n \in H_\theta(x_n, y_n)$ then $m(x_n, y_n, z_n)$ coarsely coincides with z_n, which yields
$$\lim_\omega H_\theta(x_n, y_n) \subseteq \mathrm{hull}(\boldsymbol{x}, \boldsymbol{y}).$$
To prove the other containment, suppose $\boldsymbol{z'} \in \mathrm{hull}(\boldsymbol{x}, \boldsymbol{y})$. Let $z_n = m(x_n, y_n, z'_n) \in H_\theta(x_n, y_n)$. We have $\boldsymbol{z'} = \boldsymbol{z}$ because of the definition of the median in $\boldsymbol{\mathcal{X}}$, so $\boldsymbol{z'} \in \lim_\omega H_\theta(x_n, y_n)$. □

3.2. **Topological flats in asymptotic cones.**

Lemma 3.3. *Let $\boldsymbol{\mathcal{X}}$ be an asymptotic cone of $\mathcal{X}$ and let $\mathbf{F} \subseteq \boldsymbol{\mathcal{X}}$ be a bilipschitz n–flat. Let $\boldsymbol{H}$ be an ultralimit of uniformly hierarchically quasiconvex subsets of $\mathcal{X}$ and suppose that $\mathbf{F}$ is contained in a neighborhood of $\boldsymbol{H}$ of finite radius. Then $\mathbf{F} \subseteq \boldsymbol{H}$.*

Proof. Suppose by contradiction that there exists some $p \in \mathbf{F} - \boldsymbol{H}$.

By [Bow15, Proposition 1.2, Lemma 3.3], there are arbitrarily large balls in $\mathbf{F}$ contained in a subset of $\mathbf{F}$ which is a union of blocks pairwise intersecting, if at all, in a common face. We let $\mathbf{F}'$ be such a union of blocks which contains a ball around $p \in \mathbf{F}$ of radius much larger than $\sup_{x \in \mathbf{F}} \mathsf{d}(x, \boldsymbol{H})$.

After possibly subdividing the cubulation of $\mathbf{F}'$, there is a ν–block B_0 of $\mathbf{F}'$ containing p and disjoint from $\boldsymbol{H}$. Since blocks are convex hulls of any pair of opposite corners, by Lemma 3.2, B_0 is the ultralimit $\boldsymbol{H}_0$ of hulls of pairs of points. Recall from Section 1.5 that $\mathfrak{g}_{\boldsymbol{H}_0}(\boldsymbol{H})$

is a median convex subspace, so it must be a sub-block B' of B_0. If B' has dimension i then Lemma 1.19.(3) provides an $(i+1)$–dimensional topologically embedded copy of $[0,1]^{i+1}$ in $\boldsymbol{\mathcal{X}}$. This implies $i < \nu$.

For any face B_2 of B_0 not intersecting B', there exists a block B_1' whose intersection with B_0 is B_2, so that $B_1 = B_0 \cup B_1'$ is a block by [Bow15, Lemma 3.2]. We claim $\mathfrak{g}_{B_1}(\boldsymbol{H}) = \mathfrak{g}_{B_0}(\boldsymbol{H})$, which implies that B_1 is also disjoint from $\boldsymbol{H}$.

To prove the claim, note that $B' = \mathfrak{g}_{B_0}(\boldsymbol{H}) = \mathfrak{g}_{B_0}(\mathfrak{g}_{B_1}(\boldsymbol{H}))$. Since $\mathfrak{g}_{B_0}|_{B_1}$ is just the natural retraction, which is one-to-one on B', the claim follows.

We can now proceed inductively until we find a block B_n that we cannot extend to a block B_{n+1} using the procedure above, implying that we reached the boundary of $\mathbf{F}'$. Inductively, we have $\mathfrak{g}_{B_n}(\boldsymbol{H}) = \mathfrak{g}_{B_0}(\boldsymbol{H})$, but this is impossible because there are points of $\boldsymbol{H}$ that are much closer to B_n than to B_0. This is the required contradiction. □

3.3. Quasiflats and hulls.

Proposition 3.4. *Let $F\colon \mathbb{R}^\nu \to \mathcal{X}$ be a quasiflat. Then, there exists N so that the following holds. For any $\epsilon > 0$ and every R_0 there exists a ball $B \subseteq \mathbb{R}^\nu$ of radius $R \geqslant R_0$ and a set $A \subseteq \mathcal{X}$ with $|A| \leqslant N$ so that $F(B) \subseteq \mathcal{N}_{\epsilon R}(H_\theta(A))$.*

Proof. Let $\boldsymbol{\mathcal{X}}$ be any asymptotic cone of $\mathcal{X}$ with observation points a constant sequence $(F(x_0))$. Let $\boldsymbol{F}\colon \mathbb{R}^\nu \to \boldsymbol{\mathcal{X}}$ be the corresponding ultralimit of F. Let $\boldsymbol{B}$ be a ball of radius 1 in $\mathbb{R}^\nu$. By [Bow15, Proposition 1.2], $\boldsymbol{F}(\boldsymbol{B})$ is contained in a finite union of blocks. Notice that each block is the convex hull of a pair of opposite corners. The cardinality of the number of corners provides the desired N. By Lemma 3.2, $\boldsymbol{F}(\boldsymbol{B})$ is contained in the ultralimit of hulls of pairs of points. Thus, $\boldsymbol{F}(\boldsymbol{B})$ is contained in the ultralimit of a sequence of hulls of sets of at most N points (the hull of a union contains the union of the hulls).

Now, suppose by contradiction that the conclusion of the proposition fails. Then for each N, and in particular the N we found above, there is $\epsilon > 0$ so that, for all balls $B(x_0, R)$ of sufficiently large radius R, we have that $F(B(x_0,R))$ cannot be contained in $\mathcal{N}_{\epsilon R}(H_\theta(A))$ for any $A \subseteq \mathcal{X}$ with $|A| \leqslant N$. Let $B_n = B(x_0, R_n)$, where (R_n) is the scaling factor of the chosen asymptotic cone $\boldsymbol{\mathcal{X}}$, so that $\boldsymbol{B}$ is the ultralimit of the B_n. The fact that $\boldsymbol{F}(\boldsymbol{B})$ is contained in the ultralimit of a sequence of hulls $H_\theta(A_n)$ of sets A_n of at most N points implies that, for ω–a.e. n, $F(B_n)$ is contained in $\mathcal{N}_{\epsilon R}(H_\theta(A_n))$, a contradiction. □

Proposition 3.5. *For every K, N there exist $\epsilon > 0$, R_0 and L with the following property. Let B be a ball of radius $R \geqslant R_0$ in $\mathbb{R}^\nu$, and let $F\colon B \to \mathcal{X}$ be a (K,K)–quasi-isometric embedding. Let $A \subseteq \mathcal{X}$ have $|A| \leqslant N$, and suppose that $F(B) \subseteq \mathcal{N}_{\epsilon R}(H_\theta(A))$. Then $F(B') \subseteq \mathcal{N}_L(H_\theta(A))$, where B' is the sub-ball of B with the same center and radius $R/2$.*

Proof. If not, there exist constants K, N and:

- balls $B_m = B_m(0)$ of radius R_m in $\mathbb{R}^\nu$, and (K,K)–quasi-isometric embeddings $F_m : B_m \to \mathcal{X}$,
- subsets $A_m \subseteq \mathcal{X}$ with $|A_m| \leqslant N$ and

$$\lim_{m\to\infty} \frac{1}{R_m} \sup_{x \in B_m} \mathsf{d}(F_m(x), H_\theta(A_m)) = 0,$$

  but $\lim_{m\to\infty} \sup_{x \in B_{R_m/2}(0)} \mathsf{d}(F_m(x), H_\theta(A_m)) = \infty$.

We define $\ell_m(t) = \sup_{x \in F_m(B_{\min\{t,R_m\}}(0))} \mathsf{d}(x, H_\theta(A_m))$. The ultrapower $\boldsymbol{\ell}$ of the ℓ_m can be regarded as a function $\boldsymbol{\ell}\colon {}^\omega\mathbb{R}_+ \to {}^\omega\mathbb{R}_+$. Note that $\boldsymbol{\ell}$ is non-decreasing.

Let $\sigma \in {}^\omega\mathbb{R}_+$ be represented by $\boldsymbol{R}$. For $\boldsymbol{S}, \boldsymbol{T} \in {}^\omega\mathbb{R}_+$ we write $\boldsymbol{S} \ll \boldsymbol{T}$ if $\lim_\omega S_m/T_m = 0$, and we write $\boldsymbol{S} < \infty$ if $\lim_\omega S_m \neq \infty$, i.e. if $\boldsymbol{S} \gg 1$ does not hold. We find a contradiction (with the second bullet above) provided we show $\boldsymbol{\ell}(\sigma/2) = \lim_{\omega,m} \ell_m(R_m/2) < \infty$.

The first part of the second bullet above implies that $\boldsymbol{\ell}(\sigma) \ll \sigma$. We first need:

Claim 3.6. *For $\lambda \in {}^\omega\mathbb{R}_+$, if $\boldsymbol{\ell}(\lambda) \gg 1$, then for any $\alpha \gg 1$ we have $\boldsymbol{\ell}(\lambda - \alpha\boldsymbol{\ell}(\lambda)) \ll \boldsymbol{\ell}(\lambda)$.*

Proof of Claim 3.6. Suppose not. Consider an asymptotic cone $\boldsymbol{\mathcal{X}}$ of $\mathcal{X}$ with the observation point in $F(B_{\lambda-\alpha\boldsymbol{\ell}(\lambda)}(0))$ and scaling factor $\boldsymbol{\ell}(\lambda-\alpha\boldsymbol{\ell}(\lambda))$. Then any point in the image of $\boldsymbol{F}$ has distance from $\boldsymbol{H}$ bounded above by $\boldsymbol{\ell}(\lambda)/\boldsymbol{\ell}(\lambda-\alpha\boldsymbol{\ell}(\lambda)) < \infty$. In fact, any point of the image of F which gives a point of $\boldsymbol{\mathcal{X}}$ lies in a ball of radius $\lambda-\alpha\boldsymbol{\ell}(\lambda)+t\boldsymbol{\ell}(\lambda-\alpha\boldsymbol{\ell}(\lambda)) \leqslant \lambda-\alpha\boldsymbol{\ell}(\lambda)+t\boldsymbol{\ell}(\lambda)$ for some finite t, and hence in particular in the image of the ball of radius λ.

By Lemma 3.3 we have $\boldsymbol{F} \subseteq \boldsymbol{H}$. But, we chose an arbitrary observation point in $F(B_{\lambda-\alpha\boldsymbol{\ell}(\lambda)}(0))$, and thus we get a contradiction by choosing a point that maximizes the distance from $H_\theta(A)$. □

We claim that there exists $T_0 \in \mathbb{R}_+$ so that the following holds for ω–a.e. m: if $\ell_m(t) \geqslant T_0$ for some t, and $\alpha \geqslant T_0$, then $\ell_m(t-\alpha\ell_m(t)) \leqslant \ell_m(t)/2$.

The proof follows from Claim 3.6 by an application of the principle from nonstandard analysis called *underspill*; we now provide a proof in the language of ultrafilters. For each $n \in \mathbb{N}$, let $\mathcal{U}_n$ be the set of $m \geqslant n$ for which there exists $t_{m,n}, \alpha_{m,n} \in \mathbb{R}_+$ so that $\ell_m(t_{m,n}) \geqslant n$ and $\alpha_{m,n} \geqslant n$ and $\ell_m(t_{m,n} - \alpha_{m,n}\ell_m(t_{m,n})) > \ell_m(t_{m,n})/2$. Suppose that our claim does not hold, i.e. suppose the desired T_0 does not exist. Then, for arbitrarily large n, we have that $m \in \mathcal{U}_n$ for ω–a.e. m. For each m, let $n(m)$ be the maximal n for which $m \in \mathcal{U}_n$. Our assumption, and the fact that $m \notin \mathcal{U}_n$ for $n > m$, ensures that $n(m)$ exists for ω–a.e. m.

Let $\lambda \in^\omega \mathbb{R}_+$ be the ultralimit of $t_{m,n(m)}$ and let α be that of $\alpha_{m,n(m)}$. Then $\ell(\lambda) \gg 1$ and $\alpha \gg 1$, so Claim 3.6 implies that $\boldsymbol{\ell}(\lambda-\alpha\boldsymbol{\ell}(\lambda)) \ll \boldsymbol{\ell}(\lambda)$. This contradicts that $\ell_m(t_{m,n(m)} - \alpha_{m,n(m)}\ell_m(t_{m,n(m)})) > \ell_m(t_{m,n(m)})/2$ for ω–a.e. m. Thus we have T_0 with the claimed property for ω–a.e. m.

Fix one such m, which furthermore satisfies $\ell_m(R_m) \leqslant R_m/(4\alpha_0)$ (which is satisfied by ω–a.e. m by the second bullet). Let $R_m^j = R_m(1+2^{-j})/2$. In particular, $R_m^0 = R_m$.

Claim 3.7. *Either $\ell(R_m^j) \leqslant \ell_m(R_m)/2^j$ or there exists $i \leqslant j$ with $\ell_m(R_m^i) < T_0$.*

Proof of Claim 3.7. We argue by induction on j. Suppose that R_m^j satisfy $\ell_m(R_m^j) \leqslant \ell_m(R_m)/2^j$ and $\ell_m(R_m^j) \geqslant T_0$. Note that $R_m^{j+1} = R_m^j - 2^{-j-2}R_m = R_m^j - \alpha_m^j\ell_m(R_m^j)$ for some $\alpha_m^j \geqslant T_0$. Hence, the claim gives $\ell_m(R_m^{j+1}) \leqslant \ell_m(R_m^j)/2 \leqslant \ell(R_m)/2^{j+1}$, as required. □

In either of the two cases provided by Claim 3.7, there exists j with $\ell_m(R_m^j) < T_0$. This implies $\ell_m(R_m/2) < T_0$, and hence $\boldsymbol{\ell}(\sigma/2) < T_0$, as required. □

Combining Proposition 3.4 and Proposition 3.5, one gets:

Corollary 3.8. *For every quasi-isometric embedding $f\colon \mathbb{R}^n \to \mathcal{X}$, there exist L, N so that the following holds. Then there exist arbitrarily large R so that for the ball B of radius R around 0, there is a set $A_R \subset \mathcal{X}$ with $|A_R| \leqslant N$ and $f(B') \subseteq \mathcal{N}_L(H_\theta(A_R))$, where B' is as in Proposition 3.5.*

4. Orthants and quasiflats

From now on, we fix an asymphoric HHS $(\mathcal{X}, \mathfrak{S})$ of rank ν.

4.1. Orthants in $\mathcal{X}$.

We fix once and for all D so that for any $U \in \mathfrak{S}$ any two points in F_U are connected by a D–hierarchy path.

Definition 4.1. Let $U_1, \ldots, U_k$ be a pairwise-orthogonal family and let γ_i be a D–hierarchy ray in F_{U_i} so that $\pi_{U_i}(\gamma_i)$ is unbounded. We call the image of $\gamma_1 \times \cdots \times \gamma_k \subseteq F_{U_1} \times \cdots \times F_{U_k}$ under the standard embedding a *standard k–orthant* in $\mathcal{X}$ with support set $\{U_i\}$.

A *standard orthant* is a standard ν–orthant.

Remark 4.2. Observe that if $Q = \gamma_1 \times \cdots \times \gamma_k \subseteq F_{U_1} \times \cdots \times F_{U_k}$ is a standard k–orthant, then it has uniformly bounded projection to $\mathcal{C}U$ unless $U \sqsubseteq U_i$ for some i. More precisely, each γ_i has uniformly bounded projection to $\mathcal{C}U$ unless $U \sqsubseteq U_i$ (in particular, $\pi_U(\gamma_i)$ is uniformly bounded for $U \sqsubseteq U_j, j \neq i$). For each i and each $U \sqsubseteq U_i$, we have that $\pi_U(Q)$ uniformly coarsely coincides with $\pi_U(\gamma_i)$.

Lemma 4.3. *Consider a standard k–orthant O whose support set $\{U_i\}$ has the property that, for some C, we have $\min\{\mathsf{diam}_{\mathcal{C}U}(\pi_U(O)), \mathsf{diam}_{\mathcal{C}V}(\pi_V(O))\} \leqslant C$ whenever $U, V \sqsubseteq U_i$ are orthogonal and $i \leqslant k$. Then O is κ–hierarchically quasiconvex, where κ depends on $C, D, \mathcal{X}, \mathfrak{S}$.*

In particular, there exists a function κ, depending on $(\mathcal{X}, \mathfrak{S}), D$, and the asymphoricity constant, so that standard orthants are κ–hierarchically quasiconvex, and the same holds for standard k–orthants contained in standard orthants.

Proof. Let O be a standard k–orthant which is the image of $\prod_{i=1}^k \gamma_i$, where each γ_i is a hierarchy path in F_{U_i} and $\{U_1, \ldots, U_k\}$ is a pairwise orthogonal set supporting O, and let C be the given constant.

By Remark 4.2 and the fact that hierarchy paths project close to geodesics, $\pi_U(O)$ is uniformly quasiconvex in $\mathcal{C}U$, for $U \in \mathfrak{S}$.

Suppose $x \in \mathcal{X}$ has the property that $\pi_U(x)$ lies uniformly close to $\pi_U(O)$ for each $U \in \mathfrak{S}$; to verify hierarchical quasiconvexity of O, we must bound the distance from x to O.

By hierarchical quasiconvexity of $\prod_j F_{U_j}$, our x must lie uniformly close to $\prod_j F_{U_j}$, so it suffices to show that $\mathfrak{g}_{F_j}(x)$ lies uniformly close to γ_j for each j, where F_j denotes the parallel copy of F_j containing the "corner" of O. Since $\pi_U(x)$ coarsely coincides with $\pi_U(\mathfrak{g}_{F_j}(x))$ when $U \sqsubseteq U_i$, this follows from hierarchical quasiconvexity of γ_j, i.e., Lemma 4.4. □

Lemma 4.4. *Let $\gamma \colon I \to \mathcal{X}$ be a (D, D)–hierarchy path, where $I \subseteq \mathbb{R}$ is an interval. Suppose that there exists C so that, whenever $U \perp V$, either $\pi_U(\gamma)$ or $\pi_V(\gamma)$ has diameter bounded by C. Then γ is κ–hierarchically quasiconvex, where $\kappa = \kappa(D, \mathcal{X}, \mathfrak{S}, C)$.*

Proof. Let $i, j \in I$ and let $x = \gamma(i), y = \gamma(j)$. Choose $M \geqslant \max\{C, M_0\}$, where M_0 is the constant from Theorem 2.1. By Theorem 2.1, there exists C_1, depending on M, $\mathfrak{S}$ and $\mathcal{X}$, so that there is a CAT(0) cube complex $\mathcal{C}(x, y)$ and a C_1–quasimedian (C_1, C_1)–quasi-isometric embedding $\mathcal{C}(x, y) \to \mathcal{X}$ whose image C_1–coarsely coincides with $H_\theta(x, y)$. Since $\gamma|_{[i,j]}$ is a hierarchy path from x to y, $\gamma([i, j])$ is coarsely (depending on D) contained in $H_\theta(x, y)$ and hence coarsely (depending on C_1, D) contained in the image of $\mathcal{C}(x, y)$. On the other hand, the dimension bound from Theorem 2.1, the hypothesized property of C, and our choice of $M \geqslant C$ imply that $\dim \mathcal{C}(x, y) \leqslant 1$. Moreover, Theorem 2.1 implies that $\mathcal{C}(x, y)$ is the convex hull of a set of at most two 0–cubes in $\mathcal{C}(x, y)$, so $\mathcal{C}(x, y)$ is a subdivided interval. Hence $\gamma([i, j])$ and $H_\theta(x, y)$ uniformly coarsely coincide.

Now fix ϵ and suppose $x \in \mathcal{X}$ has the property that $\pi_U(x)$ lies ϵ–close to the unparameterized (D, D)–quasigeodesic $\pi_U(\gamma)$ for each $U \in \mathfrak{S}$. Then there exists $i \geqslant 0$ so that x lies ϵ–close to the image of $\pi_U \circ \gamma|_{[0,i]}$ for all U. Hence x lies κ–close to $H_\theta(\gamma(0), \gamma(i))$, where κ depends only on ϵ and the quasiconvexity function for hulls of pairs of points. But by the above discussion, this implies that x lies uniformly close to $\gamma([0, j])$, as required. □

Next we show that suitable quasi-isometric embeddings of cubical orthants have images which are approximated by standard orthants.

Lemma 4.5. *Let O be an ν–dimensional cubical orthant with a quasimedian quasi-isometric embedding $q \colon O \to \mathcal{X}$. Then there is a standard orthant $Q \subset \mathcal{X}$ with $\mathsf{d}_{haus}(q(O), Q) < \infty$.*

Proof. Let λ be so that q is λ–quasimedian and a (λ, λ)–quasi-isometric embedding.

Related points and pairs: We say that $x, y \in O$ are *i–related*, for $1 \leqslant i \leqslant \nu$, if they only differ in the i^{th} coordinate. The i–related pairs x, y and x', y' are *j–related*, for $i \neq j$, if the pairs x, x' and y, y' are j–related (i.e. if x, x', y, y' are the vertices of a rectangle in the (i, j)–plane).

Relevant domains: Let $M = M(\lambda, \mathcal{X})$ be sufficiently large. For $1 \leqslant i \leqslant \nu$, let $\mathcal{U}_i$ be the collection of all $U \in \mathfrak{S}$ so that there exist i–related $x, y \in O$ with $\mathsf{d}_U(q(x), q(y)) \geqslant M$. For any K, we also let $\mathbf{Rel}_K(q(O)) = \{U \in \mathfrak{S} : \mathsf{diam}_{\mathcal{C}U}(\pi_U(q(O))) \geqslant K\}$.

We now prove two claims about i–related pairs and $\cup_i \mathcal{U}_i$:

Claim 4.6. *There exists $C = C(\lambda, \mathcal{X})$ so that the following holds. Suppose that the i–related pairs x, y and x', y' are j–related. Then for any $U \in \mathfrak{S}$ either*

- $\mathsf{d}_U(x, y) \leqslant C$ *and* $\mathsf{d}_U(x', y') \leqslant C$*, or*
- $\mathsf{d}_U(x, x') \leqslant C$ *and* $\mathsf{d}_U(y, y') \leqslant C$*.*

Proof of Claim 4.6. Let $m : O^3 \to O$ be the median on O coming from the cubical structure (so each cube is an ℓ_1 ν–cube of unit side length). We have $m(x', x, y) = x$, so that in each $U \in \mathfrak{S}$ we have that $\pi_U(x)$ lies uniformly close to geodesics $[\pi_U(x'), \pi_U(y)]$. Similarly, $\pi_U(y')$ lies uniformly close to geodesics $[\pi_U(x'), \pi_U(y)]$. Also, $\pi_U(x')$ and $\pi_U(y)$ lie uniformly close to geodesics $[\pi_U(x), \pi_U(y')]$, forcing the endpoints of $[\pi_U(x'), \pi_U(y)]$ and $[\pi_U(x), \pi_U(y')]$ to be uniformly close in pairs, as required. $\square$

Claim 4.7. *For M sufficiently large, $U \perp V$ whenever $U \in \mathcal{U}_i, V \in \mathcal{U}_j$ and $i \neq j$.*

Proof of Claim 4.7. Consider distinct i, j, an i–related pair x, y and some U with $\mathsf{d}_U(q(x), q(y)) \geqslant M$, and a j–related pair w, z and some V so that $\mathsf{d}_V(q(w), q(z)) \geqslant M$.

Provided $M \geqslant 10(\nu - 1)C$, applying Claim 4.6 at most $\nu - 1$ times allows us to change the coordinates of w, z (other than the j^{th}) to find an i–related pair x', y' which is j–related to x, y. Moreover, we have:

- $\mathsf{d}_V(q(x), q(x')) \geqslant M/2$ and $\mathsf{d}_V(q(y), q(y')) \geqslant M/2$;
- $\mathsf{d}_U(q(x), q(y)) \geqslant M$ and $\mathsf{d}_U(q(x'), q(y')) \geqslant M/2$.

Claim 4.6 implies that $\mathsf{d}_U(q(x), q(x')) \leqslant C$, $\mathsf{d}_U(q(y), q(y')) \leqslant C$ and $\mathsf{d}_V(q(x), q(y)) \leqslant C$, $\mathsf{d}_V(q(x'), q(y')) \leqslant C$.

For M large enough, this implies that $U \perp V$. Indeed, if $U = V$, then the triangle inequality yields $4C \geqslant M/2$, a contradiction. If $U \pitchfork V$, then there exists $p \in \{x, x', y, y'\}$ with $\pi_U(p)$ E–far from ρ^V_U and $\pi_V(p)$ E–far from ρ^U_V, contradicting consistency. A similar contradiction arises if U, V are $\subsetneq$–comparable. Hence $U \perp V$, as required. $\square$

The candidate standard orthant: Let γ'_i be the image of the axis along the i^{th} coordinate in O. Since q is quasimedian and a quasi-isometric embedding, γ'_i is a quasi-geodesic projecting to unparameterized quasi-geodesics in every $\mathcal{C}U$, i.e. it is a $D' = D'(\lambda)$–hierarchy ray. By [DHS17, Lemma 3.3], there exist $U^i_1, \ldots, U^i_{k_i}$ so that $\pi_{U^i_j}(\gamma'_i)$ is unbounded.

For $1 \leqslant i \leqslant \nu,\ 1 \leqslant j < j' \leqslant k_i$, we have $U^j_i \perp U^{j'}_i$. Since each $U^i_j \in \mathcal{U}_i$, Claim 4.7 and the fact that $\mathcal{X}$ has rank ν implies that $k_i = 1$ for each i. To streamline notation, let $U_i = U^i_1$.

Since $\{U_1, \ldots, U_\nu\}$ is a pairwise-orthogonal set, the following holds for all $i \leqslant \nu$: if $U, V \subsetneq U_i$ have $\mathsf{diam}(\mathcal{C}U), \mathsf{diam}(\mathcal{C}V) > E$, then $U \not\perp V$, for otherwise $\{U_1, \ldots, U_{i-1}, U, V, U_{i+1}, \ldots, U_\nu\}$ would contradict that $\mathcal{X}$ is asymphoric. It follows from Corollary 2.16 that each F_{U_i} is hyperbolic. Hence there exists a D''–hierarchy ray γ_i in F_{U_i} so that the distance between $\gamma_i(t)$ and $\gamma'_i(t)$ is uniformly bounded for all $t \in [0, \infty)$.

The γ_i define a standard orthant Q with support $\{U_i\}$.

$q(O)$ **and Q lie within finite Hausdorff distance:** We claim the following. For $p \in O$ we denote by p_i the point on the i–th coordinate axis with the same i–th coordinate as p.

Then there exists C' so that $d_{\mathcal{C}U}(q(p), q(p_i)) \leqslant C'$ whenever $U \notin \bigcup_{j\neq i} \mathcal{U}_j$. This holds because we can find a sequence of at most ν points, starting with p and ending with p_i, so that consecutive elements are j–related for $j \neq i$. By definition, if consecutive elements have far away projection to some $\mathcal{C}U$, then $U \in \mathcal{U}_j$ for $j \neq i$.

Let now $p \in O$. By the above claim, $\pi_U(q(p))$ coarsely coincides with $\pi_U(q(p_i))$ if $U \in \mathcal{U}_i$, and otherwise it coarsely coincides with $\pi_U(c)$, where c is the image of the "corner" of O. We can find points $\gamma_i(t_i)$ uniformly close to $q(p_i) \in \gamma_i'$, and the $\gamma_i(t_i)$ define a point p' of Q. It is readily checked that for every U, $\pi_U(q(p))$ coarsely coincides with $\pi_U(p')$, so that $q(p)$ and p' are within uniformly bounded distance. This proves that $q(O)$ is contained in a finite radius neighborhood of Q. A very similar argument proves the other containment. □

4.2. **Coarse intersections of orthants.**

Definition 4.8. Let $A, B \subset \mathcal{X}$. Suppose that there exists R_0 so that for any $R, R' \geqslant R_0$, we have $\mathsf{d}_{haus}(\mathcal{N}_R(A) \cap \mathcal{N}_R(B), \mathcal{N}_{R'}(A) \cap \mathcal{N}_{R'}(B)) < \infty$. Then we refer to any subspace at finite Hausdorff distance from $\mathcal{N}_{R_0}(A) \cap \mathcal{N}_{R_0}(B)$ as the *coarse intersection of A and B*, which we denote $A \tilde{\cap} B$.

Lemma 4.9. *Let A, B be hierarchically quasiconvex. Then $A \tilde{\cap} B$ is well-defined and coarsely coincides with $\mathfrak{g}_A(B)$.*

Proof. It is easily seen from Lemma 1.19.(3) that any point in $\mathfrak{g}_A(B)$ lies within uniformly bounded distance of both A and B. On the other hand, if $p \in \mathcal{N}_R(A) \cap \mathcal{N}_R(B)$ then p is close to $\mathfrak{g}_A(p')$ for some $p' \in B$ which is R–close to p. □

Lemma 4.10. *Let O, O' be standard orthants in $\mathcal{X}$ with supports $\{U_i\}_{i\leqslant\nu}, \{U_i'\}_{i\leqslant\nu}$. Then $O \tilde{\cap} O'$ is well-defined, and coarsely coincides with $\mathfrak{g}_O(O')$, as well as with a standard k–orthant whose support is contained in $\{U_i\}_{i\leqslant\nu} \cap \{U_i'\}_{i\leqslant\nu}$.*

Proof. By Lemma 4.9, we only need to show that $\mathfrak{g}_O(O')$ coarsely coincides with a standard k–orthant whose support is contained in $\{U_i\} \cap \{U_i'\}$.

Let γ_i be the hierarchy ray in F_{U_i} participating in O, and similarly for γ_i' and O'. Let $\{V_j\}_{j=1,\dots,k}$ be the set of all $V_j = U_i = U_{i'}'$ so that γ_i and $\gamma_{i'}'$ lie within bounded Hausdorff distance, in which case set $\alpha_j = \gamma_i$. Let O'' be a standard k–orthant contained in O with support set $\{V_j\}$ defined by the α_j. We claim that O'' represents $O \tilde{\cap} O'$.

By Lemma 4.3, O'' is hierarchically quasiconvex, and $G = \mathfrak{g}_O(O')$ is hierarchically quasiconvex by Lemma 1.19.(1). We claim that O'' coarsely coincides with G. Since they are hierarchically quasiconvex, we only need to argue that their projections to each $\mathcal{C}U$ coarsely coincide.

By Remark 4.2, for each U, $\pi_U(O'')$ coarsely coincides with some $\pi_U(\alpha_j)$. In particular, if U is not nested in some U_j, then $\pi_U(O'')$ uniformly coarsely coincides with with each $\pi_U(\alpha_j(0))$. Also, $\pi_U(G)$ coarsely coincides with the projection of a single γ_i, if $\gamma_i = \alpha_j$ for some j. Otherwise $\pi_U(G)$ coarsely coincides with $\pi_U(\alpha_j(0))$ for each j. Hence $\pi_U(G)$ and $\pi_U(O'')$ coarsely coincide for all U. □

4.3. **Quasiflats theorem.**

Theorem 4.11. *Let $\mathcal{X}$ be an asymphoric HHS of rank ν and let $f \colon \mathbb{R}^\nu \to \mathcal{X}$ be a quasi-isometric embedding. Then there exist finitely many standard orthants $Q_i \subseteq \mathcal{X}$ for $1 \leqslant i \leqslant k$, so that*

$$\mathsf{d}_{haus}(f(\mathbb{R}^\nu), \cup_{i=1}^k Q_i) < \infty.$$

Proof. Let L, N be as in Corollary 3.8. Then there exist an increasing unbounded sequence $R_1 < R_2 < \dots$ and sets $A_i \subseteq \mathcal{X}$ of cardinality at most N so that the following holds. Let B_i be the ball in $\mathbb{R}^\nu$ of radius R_i centered at a fixed basepoint, and let $H_i = H_\theta(A_i)$. Then

$f(B_i) \subseteq \mathcal{N}_L(H_i)$. Let $c_i\colon \mathcal{Y}_i \to H_i$ be the (C,C)–quasi-isometry provided by Theorem 2.1, so $\mathcal{Y}_i$ is a CAT(0) cube complex of dimension $\leqslant \nu$; the constant C depends on N.

Now we pass to (non-rescaled!) ultralimits[1]. More specifically, f has an ultralimit which is a (K,K)–quasi-isometric embedding $\hat{f}\colon \mathbb{R}^\nu \to \widehat{\mathcal{X}}$, for some ultralimit $\widehat{\mathcal{X}}$ of $\mathcal{X}$. It is easily deduced from Corollary 2.15 that $\widehat{\mathcal{X}}$ is a coarse median space and we have the following: there is a CAT(0) cube complex $\hat{\mathcal{Y}}$, an ultralimit of the $\mathcal{Y}_i$, endowed with a C–quasimedian (C,C)–quasi-isometry $\hat{c} : \hat{\mathcal{Y}} \to \widehat{\mathcal{X}}$ so that the image of $\hat{f}$ lies in the L–neighborhood of $\mathrm{im}(\hat{c})$.

By Theorem 1.1 of [Hua14b], there exist n–dimensional cubical orthants $O_1,\dots,O_k$ in $\hat{\mathcal{Y}}$ so that $\mathsf{d}_{haus}(\hat{f}(\mathbb{R}^\nu), \hat{c}(\cup_{i=1}^k O_i)) < \infty$. Moreover, $\hat{c}(O_i)$ lies within finite Hausdorff distance of $\hat{f}(O_i')$ for some $O_i' \subseteq \mathbb{R}^\nu$. Hence, $Q_i = f(O_i')$ is the image of a C'–quasimedian (C',C')–quasi-isometric embedding, and hence by Lemma 4.5 it lies within finite Hausdorff distance of a standard orthant. The Q_i are as required. □

4.4. **Controlled number of orthants.** We now improve Theorem 4.11, by showing that the number of standard orthants required can be bounded in terms of the quasi-isometry constants:

Theorem 4.12. *Let $\mathcal{X}$ be an asymphoric HHS of rank ν. For every K there exists N so that the following holds. Let $f\colon \mathbb{R}^\nu \to \mathcal{X}$ be a (K,K)–quasi-isometric embedding. Then there exist standard orthants $Q_i \subseteq \mathcal{X}$, $i = 1,\dots,N$, so that $\mathsf{d}_{haus}(f(\mathbb{R}^\nu), \cup_{i=1}^N Q_i) < \infty$.*

The following is a slightly stronger version of the well-known fact that quasi-isometric embeddings of $\mathbb{R}^n$ into itself are coarsely surjective, see [KL97a, Corollary 2.6].

Lemma 4.13. *For every $K, n \geqslant 1$ there exists C so that the following holds. Let $f\colon \mathbb{R}^n \to \mathbb{R}^n$ be a (K,K)–coarsely Lipschitz proper map. Then $\mathsf{d}_{haus}(f(\mathbb{R}^n), \mathbb{R}^n) \leqslant C$.*

Proof. We actually show that if $f\colon \mathbb{R}^n \to \mathbb{R}^n$ is continuous and proper, then f is surjective, and the lemma follows from the fact that f can be approximated by a continuous map.

Since f is proper, it extends to a continuous map $\overline{f} : \overline{\mathbb{R}^n} \to \overline{\mathbb{R}^n}$ between two copies of the 1–point compactification $\overline{\mathbb{R}^n}$ of $\mathbb{R}^n$, which is homeomorphic to the sphere S^n. Also, it is easily seen that we can identify the domain $\overline{\mathbb{R}^n}$ with S^n in such a way that, since f is coarsely Lipschitz, no pair of antipodal points have the same image. But then $\overline{f}$ must be surjective, for otherwise the Borsuk-Ulam theorem would force the existence of such pair of antipodal points. Since $\overline{f}$ is surjective, then so is f, as required. □

Proposition 4.14. *For every K there exists N so that the following holds. Let $F\colon \mathbb{R}^\nu \to \mathcal{X}$ be a (K,K)–quasi-isometric embedding whose image lies at finite Hausdorff distance from $\bigcup_{i=1}^k O_i$, where each O_i is a standard orthant. If $\mathsf{d}_{haus}(O_i, O_j) = \infty$ when $i \neq j$, then $k \leqslant N$.*

Proof. The idea of the proof is that each of the k orthants contributes at least ϵR^ν volume growth to $F(\mathbb{R}^\nu)$, but the volume growth of $F(\mathbb{R}^\nu)$ is bounded above by R^ν times a (large) constant depending on K.

Let $D = \mathsf{d}_{haus}(F(\mathbb{R}^\nu), \bigcup_{i=1}^k O_i)$. By Lemma 4.10, since the O_i are pairwise at infinite Hausdorff distance, for each i we can find a sub-orthant $O_i' \subset O_i$ so that for each i,j, $\mathsf{d}(O_i', O_j') \geqslant 2D + 1$. We will identify O_i' with $[0,\infty)^\nu$.

Let $A_i \subseteq \mathbb{R}^\nu$ be the set of points whose image under F is at distance at most D from O_i'. Note that the A_i are disjoint. For each R and i, there exists a sub-orthant $O_i^R \subset O_i'$ so that if $x \in A_i$ and $\mathsf{d}(F(x), O_i^R) \leqslant D$, then $B_R(x) \subseteq A_i$.

Let g_i be the composition of F and the gate map to O_i'; the map g_i is (K',K')–coarsely Lipschitz for some $K' = K'(K,\mathcal{X})$. Let C be as in Lemma 4.13 for K'.

[1] If $\mathcal{X}$ is proper, one can take Hausdorff limits instead. To avoid that assumption, we use ultralimits instead. If $\mathcal{X}$ is not proper then $\widehat{\mathcal{X}}$ is (much) bigger than $\mathcal{X}$.

We claim that there are suborthants $O_i'' \subset O_i'$ so that $O_i'' \subset N_C(g_i(A_i)) \cap O_i'$. If not, for each n there exist $p_n \in A_i$ with $g_i(p_n) \in O_i^n$ and some $x_n \in O_i^n$ with $\mathsf{d}(x_n, g_i(p_n)) \leqslant 2C$ but $\mathsf{d}(x_n, g_i(A_i)) > C$. Then, we consider the (non-rescaled!) ultralimit X of $\mathbb{R}^\nu$ with observation point (p_n), which is isometric to $\mathbb{R}^\nu$. This yields a (K', K')–coarsely Lipschitz map from X to an ultralimit of the O_i^n, which is again a copy of $\mathbb{R}^\nu$, but the map is not C–coarsely surjective, contradicting Lemma 4.13 and thus verifying the claim.

We now bound from below $\beta_R = |\{x \in \mathbb{Z}^\nu : F(x) \in B_R(F(0))\}|$. There exists $t = t(K)$ so that $\beta_R \leqslant tR^n$. Let $C' = C'(C, n, K)$ satisfy $O_i'' \subset N_{C'}(g_i(A_i \cap \mathbb{Z}^\nu)) \cap O_i'$. Consider a maximal $(2C' + 1)$–net N_i in O_i'' and, for any point p of the net, choose some $q \in A_i \cap \mathbb{Z}^\nu$ with $\mathsf{d}(p, F(q)) \leqslant C'$. Distinct p yield distinct q. Moreover, $|N_i| \cap B_R(F(0)) \geqslant t'R^n$ for all sufficiently large R and some $t' = t'(C', \mathcal{X})$. Since the A_i are disjoint, we have $\beta_R \geqslant kt'R^n$ for all sufficiently large R. Hence $k \leqslant t/t'$, and we are done. □

Proof of Theorem 4.12. By Theorem 4.11, the image of F lies at finite Hausdorff distance from a union of orthants $\bigcup_{i=1}^k O_i$. We can assume that $\mathsf{d}_{haus}(O_i, O_j) = \infty$ when $i \neq j$; indeed, if not, then we can drop O_i or O_j from the collection without affecting the conclusion. Hence, $k \leqslant N$, for N as in Proposition 4.14. □

4.5. **Controlled distance.** As in the cubical case, it is not possible in general to give an effective bound on the Hausdorff distance between a quasiflat and the corresponding union of orthants. However, we have the following:

Lemma 4.15. *For every K, N there exists L so that the following holds. Let $F\colon \mathbb{R}^\nu \to \mathcal{X}$ be a (K, K)–quasi-isometric embedding whose image lies at finite Hausdorff distance from $\bigcup_{i=1}^N O_i$, where each O_i is a standard orthant. Then $F \subset \mathcal{N}_L(H_\theta(\bigcup_{i=1}^N O_i))$.*

Proof. Let F and O_i be as in the statement. Any bounded set in O_i lies in a uniform neighborhood of the hull of the "corner point" of O_i and some point along the diagonal. Hence, there exists D so that any ball B in $\mathbb{R}^n$ has the property that $F(B)$ is contained in the D–neighborhood of $H_\theta(A)$ for some $A \subseteq \bigcup_i O_i$ with $|A| \leqslant 2N$. For L as in Proposition 3.5, there exist arbitrarily large balls B' in $\mathbb{R}^\nu$ so that $F(B') \subseteq \mathcal{N}_L(H_\theta(A)) \subseteq \mathcal{N}_L(H_\theta(\bigcup_{i=1}^N O_i))$ for some $A \subseteq \bigcup_i O_i$. Hence, the same holds for $\mathbb{R}^\nu$, as required. □

Corollary 4.16. *For each K there exists L, N so that the following holds. Let $F\colon \mathbb{R}^\nu \to \mathcal{X}$ be a (K, K)–quasi-isometric embedding. Then there exist standard orthants $O_1, \ldots, O_N$ so that $F \subset \mathcal{N}_L(H_\theta(\bigcup_{i=1}^N O_i))$.*

Proof. Follows immediately from Theorem 4.12 and Lemma 4.15. □

5. Induced maps on hinges: mapping class group rigidity

Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. We will impose three additional assumptions on $(\mathcal{X}, \mathfrak{S})$, which are satisfied by the standard HHS structure on the mapping class group, described in [BHS15, Section 11]. First, we introduce a few relevant definitions.

Definition 5.1 (Standard flat)**.** Let $U_1, \ldots, U_k$ be a pairwise-orthogonal family and let γ_i be a bi-infinite D–hierarchy path in F_{U_i} with $\pi_{U_i}(\gamma_i)$ unbounded. We call the image of $\gamma_1 \times \cdots \times \gamma_k \subseteq F_{U_1} \times \cdots \times F_{U_k}$ under the standard embedding a *standard k–flat* in $\mathcal{X}$ with support set $\{U_i\}$. For brevity, we refer to a standard ν–flat as a *standard flat.*

The next definition describes those subsets of $\mathfrak{S}$ which give rise to standard flats.

Definition 5.2 (Complete support set)**.** A *complete support set* is a subset $\{U_i\}_{i=1}^\nu \subset \mathfrak{S}$ whose elements are pairwise orthogonal and satisfy $\mathsf{diam}(\mathcal{C}U_i) = \infty$ for all $i \leqslant \nu$.

Note that a complete support set $\{U_i\}$ and a pair of distinct points $\{p_i^\pm\} \in \partial \mathcal{C}U_i$ for each i, allows one to construct a standard flat, $\mathcal{F}_{\{(U_i, p_i^\pm)\}}$ associated to some choice of bi-infinite hierarchy paths in each F_{U_i} whose projection to $\mathcal{C}U_i$ has limit points $\{p_i^\pm\}$ in $\mathcal{C}U_i$. Accordingly, it is easy to verify that a complete support set is the support set of some standard flat if and only if each $\partial\mathcal{C}U_i$ contains at least two points.

Definition 5.3 (Hinge, orthogonal hinges)**.** A *hinge* is a pair (U, p) with:

- $U \in \mathfrak{S}$;
- U belongs to some complete support set; and,
- $p \in \partial\mathcal{C}U$.

Let $\mathbf{Hinge}(\mathfrak{S})$ be the set of hinges. We say $(U, p), (V, q) \in \mathbf{Hinge}(\mathfrak{S})$ are *orthogonal* if $U \perp V$.

Definition 5.4 (Ray associated to a hinge)**.** A *μ–ray associated to* a hinge $\sigma = (U, p)$ is a μ–hierarchy path $\mathfrak{h}_\sigma$ so that $\pi_U(\mathfrak{h}_\sigma)$ is a quasigeodesic ray representing p and so that $\mathsf{diam}(\pi_V(\mathfrak{h}_\sigma)) \leqslant \mu$ for $V \neq U$.

Remark 5.5. Any two candidates for $\mathfrak{h}_\sigma$ lie at finite Hausdorff distance, so for our purposes an arbitrary choice is fine. If $\sigma \neq \sigma' \in \mathbf{Hinge}(\mathfrak{S})$, then $\mathsf{d}_{haus}(\mathfrak{h}_\sigma, \mathfrak{h}_{\sigma'}) = \infty$.

Remark 5.6. Each hinge corresponds to a 0–simplex in the HHS boundary $\partial\mathcal{X}$; see [DHS17].

The first additional assumption holds, for example, in any hierarchically hyperbolic group:

Assumption 1. *For every $U \in \mathfrak{S}$, either $\mathsf{diam}(\mathcal{C}U) \leqslant E$ or $|\partial\mathcal{C}U| \geqslant 2$ has at least two points at infinity.*

Remark 5.7. In what follows, we could replace Assumption 1 with: for each $U \in \mathfrak{S}$ which is the first coordinate of some hinge, $|\partial\mathcal{C}U| \geqslant 2$. Equivalently, each $U \in \mathfrak{S}$ which is the first coordinate of some hinge is the first coordinate of at least two hinges.

The second assumption roughly says that, if a standard 1–flat is contained in some standard flat, then it can be realized as the intersection of a pair of standard flats.

Assumption 2. *For every U contained in a complete support set there exist complete support sets $\mathcal{U}_1, \mathcal{U}_2$ with $\{U\} = \mathcal{U}_1 \cap \mathcal{U}_2$.*

The third assumption is a two-dimensional version of the second one; this assumption says that if a standard 2–flat is contained in a standard flat, then it can be obtained as the intersection of some pair of standard flats.

Assumption 3. *If $\nu > 2$, then for every U, V, with each contained in a complete support set and with $U \perp V$, there exist complete support sets $\mathcal{U}_1, \mathcal{U}_2$ with $\{U, V\} = \mathcal{U}_1 \cap \mathcal{U}_2$.*

Theorem 5.8. *Let $(\mathcal{X}, \mathfrak{S})$, $(\mathcal{Y}, \mathfrak{T})$ be asymphoric HHS satisfying assumptions (1), (2) and (3). For any quasi-isometry $f\colon \mathcal{X} \to \mathcal{Y}$, there exists a bijection $f^\sharp\colon \mathbf{Hinge}(\mathfrak{S}) \to \mathbf{Hinge}(\mathfrak{T})$ satisfying:*

- *$f^\sharp$ preserves orthogonality of hinges;*
- *for all $\sigma \in \mathbf{Hinge}(\mathfrak{S})$, we have $\mathsf{d}_{haus}(\mathfrak{h}_{f^\sharp(\sigma)}, f(\mathfrak{h}_\sigma)) < \infty$.*

Remark 5.9. Under suitable conditions, we expect that there exists an analogue of Theorem 5.8 in which hinges are replaced by sets of pairs $\{(U_i, p_i)\}$, where $\{U_i\}_i$ is a pairwise orthogonal set and $p_i \in \partial\mathcal{C}U_i$. In particular, one should be able to show in this way that isolated flats are taken close to isolated flats. More strongly, one could consider the situation where flats coarsely intersect in subspaces of codimension $\geqslant 2$, as in [FLS15].

Proof of Theorem 5.8. Let $\sigma = (U, p) \in \mathbf{Hinge}(\mathfrak{S})$.

How we will define $f^\sharp$: We will produce a hinge σ' so that $\mathsf{d}_{haus}(\mathfrak{h}_{\sigma'}, f(\mathfrak{h}_\sigma)) < \infty$. Remark 5.5 implies that σ' is uniquely determined by this property, so we can set $f^\sharp(\sigma) = \sigma'$. To see that this is a bijection, let $\bar{f} : \mathcal{Y} \to \mathcal{X}$ be a quasi-inverse of f. Then $\mathsf{d}_{haus}(\bar{f}(h_{\sigma'}), h_\sigma) < \infty$, so we can define an inverse for $f^\sharp$ in the same way.

Choosing σ': Since (U, p) is a hinge, U is in a complete support set $\{U_i\}_i$. By Assumption 1, $|\partial \mathcal{C} U_i| \geqslant 2$ for each i. Hence there exists a standard flat $\mathcal{F}$ with support $\{U_i\}$.

Assumption 2 provides two standard flats $\mathcal{F}_1, \mathcal{F}_2$, the intersection of whose support sets is $\{U\}$; moreover, in view of Lemma 4.10, we can arrange that $\mathcal{F}_1 \tilde{\cap} \mathcal{F}_2$ is a line coarsely containing $\mathfrak{h}_\sigma$. By Theorem 4.11, $f(\mathcal{F}_1)$ and $f(\mathcal{F}_2)$ are coarsely equal to unions of finitely many standard orthants. Hence $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ has the following three properties:

- $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is a finite union of coarse intersections of pairs of standard orthants;
- $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is coarsely $\mathbb{R}$;
- $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ coarsely contains $f(\mathfrak{h}_\sigma)$.

By Lemma 4.10 and the first of the above properties, $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is the finite union of standard k–orthants (arising as coarse intersections of pairs of standard orthants). Hence, one of these pairs gives a 1–orthant (in particular, a copy of $\mathbb{R}_+$) which coarsely coincides with $f(\mathfrak{h}_\sigma)$.

Let σ' be the hinge (V, q), where V is the domain of the orthant just determined and q is the unique point in ∂V determined by the fact that $f(\mathfrak{h}_\sigma)$ projects to a quasi-geodesic ray in $\mathcal{C}V$. Then σ' is the hinge uniquely determined by $f(\mathfrak{h}_\sigma)$, as required.

Preservation of orthogonality: Let σ, σ' be orthogonal hinges. Assumption 3 provides a standard 2–flat, $\mathcal{F}$, coarsely containing $\mathfrak{h}_\sigma$ and $\mathfrak{h}_{\sigma'}$. Moreover, $\mathcal{F}$ coarsely coincides with $\mathcal{F}_1 \tilde{\cap} \mathcal{F}_2$, for standard flats $\mathcal{F}_1, \mathcal{F}_2$.

Hence $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is a 2–dimensional quasiflat. On the other hand, by Theorem 4.11, $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is the union of finitely many coarse intersections of pairs of standard orthants, so by Lemma 4.10, $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is coarsely the union of disjoint standard 2–orthants $O_0, \ldots, O_{t-1}$. Moreover, $\mathfrak{h}_{f^\sharp(\sigma)}$ and $\mathfrak{h}_{f^\sharp(\sigma')}$ coarsely coincide with coordinate rays of some O_i, O_j.

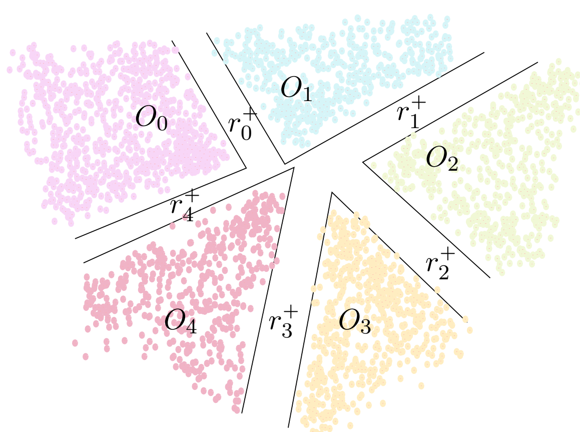

FIGURE 3. The 2–orthants $O_0, \ldots, O_t$ and the cyclic ordering of their coordinate rays (up to coarse coincidence).

Now, as shown in Figure 3, we can cyclically order the coordinate rays in $O_0, \ldots, O_{t-1}$. First, label the orthants so that for each $s \in \mathbb{Z}_t$, the 2–orthant O_s has the property that one of its coordinate rays r_s^- coarsely coincides with a coordinate ray in O_{s-1} and the other, r_s^+, coarsely coincides with a coordinate ray in O_{s+1}. Now cyclically order the coarse equivalence classes of rays: $r_0^+, r_1^+, \ldots, r_{t-1}^+$.

We claim that $\mathfrak{h}_{f^\sharp(\sigma)}$, $\mathfrak{h}_{f^\sharp(\sigma')}$ must be adjacent in this order. This will imply that they are coarsely contained in a common 2–orthant, and hence $f^\sharp(\sigma)\perp f^\sharp(\sigma')$, as required.

Indeed, if there was a coordinate ray r between $\mathfrak{h}_{f^\sharp(\sigma)}$ and $\mathfrak{h}_{f^\sharp(\sigma')}$, then r is coarsely $\mathfrak{h}_{f^\sharp(\sigma'')}$, so that by definition $f^{-1}(r)$ is coarsely $\mathfrak{h}_{\sigma''}$. (Here we used Assumption 2, which guarantees that r is the ray associated to some hinge.) But then $\mathfrak{h}_\sigma, \mathfrak{h}_{\sigma'}, \mathfrak{h}_{\sigma''}$ pairwise have infinite Hausdorff distance, are contained in the same standard 2–orthant, and they each arise as the coarse intersection with some other orthant, contradicting Lemma 4.10. □

5.1. **Sharpening of $f^\sharp$.** The hinge $f^\sharp(\sigma)$ prescribes a hierarchy ray which lies within finite distance of $f(\mathfrak{h}_\sigma)$, but it does not (and cannot) provide a uniform bound on the distance; which is what one typically needs to show that two given quasi-isometries coarsely coincide. Under many circumstances, finiteness can actually be promoted to a uniform bound, with little extra work. As an illustration of this, we give an example tailored to the mapping class group case in the following lemma.

Lemma 5.10. *Let $(\mathcal{X}, \mathfrak{S}), (\mathcal{Y}, \mathfrak{T})$ be asymphoric HHS satisfying Assumptions (1), (2) and (3). There exists C with the following property. Let $\{U_i\}_{i=1}^n \subseteq \mathfrak{S}$ be a complete support set, and let $p_i^\pm$ be distinct points in $\partial\mathcal{C}U_i$. Suppose that there exists a complete support set $\{V_i\}_{i=1}^n \subseteq \mathfrak{T}$ and distinct points $q_i^\pm \in \partial\mathcal{C}V_i$ so that for each $k = 1, \ldots, n$ we have $f^\sharp(U_k, p_k^\pm) = (V_k, q_k^\pm)$. Then, $\mathsf{d}_{haus}(f(\mathcal{F}_{\{(U_i,p_i^\pm)\}}), \mathcal{F}_{\{(V_j,q_j^\pm)\}}) \leqslant C$.*

Proof. Hierarchical quasiconvexity of $\mathcal{F}_{\{(V_j,q_j^\pm)\}}$ implies it uniformly coarsely coincides with $H_\theta(\mathcal{F}_{\{(V_j,q_j^\pm)\}})$. Containment of $f(\mathcal{F}_{\{(U_i,p_i^\pm)\}})$ in a uniform neighborhood of $\mathcal{F}_{\{(V_j,q_j^\pm)\}}$ then follows from Lemma 4.15. The other containment follows by applying the same argument to a quasi-inverse of f. □

5.2. **Mapping class groups.** We now use Theorem 5.8 to provide a new proof of quasi-isometric rigidity of mapping class groups.

Theorem 5.11. [BKMM12] *Let $\mathcal{X}$ be the the mapping class group of a non-sporadic surface S. Then for any K there exists L so that: for each quasi-isometry $f\colon \mathcal{X} \to \mathcal{X}$ there exists a mapping class g so that f L–coarsely coincides with left-multiplication by g.*

Proof. Consider the standard HHS structure on $\mathcal{X}$, so that $\mathfrak{S}$ is the collection of all essential subsurfaces, and the $\mathcal{C}Y$ are curve complexes. (For details on the structure, see [BHS15, Section 11].)

A subsurface Y lies in a complete support set if and only if it is an annulus, a once-punctured torus or a 4-holed sphere. The assumptions of Theorem 5.8 are clearly satisfied.

Consider any quasi-isometry $f\colon \mathcal{X} \to \mathcal{X}$. A hinge (U, p) is *annular* if U is an annulus. We now show that if a hinge σ is annular, then so is $f^\sharp(\sigma)$. Indeed, a hinge σ being annular is characterized by the following property: σ is contained in a maximal collection $\mathcal{H}$ of pairwise orthogonal hinges, and there exists a unique hinge σ' so that $(\mathcal{H} - \{\sigma\}) \cup \{\sigma'\}$ is a maximal pairwise orthogonal set of hinges. This property is illustrated in Figure 4, where, if σ is (U, p^+), then σ' is (U, p^-), where $\partial\mathcal{C}U = \{p^\pm\}$.

Since the bijection $f^\sharp$ preserves orthogonality and non-orthogonality, it preserves the above property, so $f^\sharp$ preserves being annular.

Note that for any annulus U, the set $\partial\mathcal{C}U$ has exactly two points. We now claim that for each annulus U there exists an annulus V so that, denoting $\{p^\pm\} = \partial\mathcal{C}U$, we have $f^\sharp(U, p^\pm) = (V, q^\pm)$ for $q^\pm \in \partial\mathcal{C}V$. This holds as above, since (U, p^-) is the only hinge that can replace (U, p^+) in a certain maximal set of pairwise orthogonal hinges (one in which the

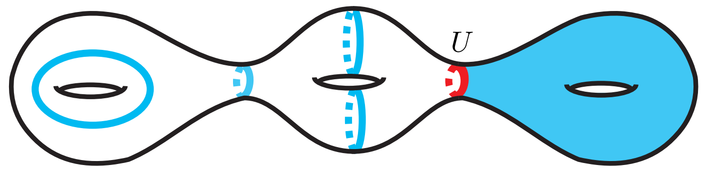

FIGURE 4. This figure shows a complete support set, consisting of five annuli and one once-punctured torus. This is the only complete support set containing all the subsurfaces except the annulus about the boundary of the once-puncture torus, denoted U in the figure. Hence U is non-replaceable.

core curve of U is a non-replaceable curve). We write $V = f^*(U)$. Notice that Lemma 5.10 now applies to show that any Dehn twist flat of $\mathcal{X}$ is mapped within uniformly bounded distance of a Dehn twist flat.

Moreover, we have a well defined simplicial automorphism ϕ of the curve graph $\mathcal{C}S$, where $\phi(\alpha) = \beta$ if $B = f^*(A)$, where the annuli A, B have core curves α, β respectively. By a theorem of Ivanov [Iva97], any simplicial automorphism of $\mathcal{C}S$ is induced by an element of the mapping class group; we denote by g the element corresponding to ϕ.

Suppose we are given a Dehn twist flat $\mathcal{F}$ with complete support set $\mathcal{U}$. Then, as noted above, $f(\mathcal{F})$ is coarsely a Dehn twist flat with complete support set $\{f^*(U)\}_{U\in\mathcal{U}} = \{gU\}_{U\in\mathcal{U}}$.

We can now conclude that for any Dehn twist flat $\mathcal{F}$, we have that $f(\mathcal{F})$ and $g\mathcal{F}$ are within bounded Hausdorff distance. For any point $x \in \mathcal{X}$, we can find Dehn twist flats $\mathcal{F}_1^x, \mathcal{F}_2^x$ that have neighborhoods of uniformly bounded radius whose intersection contains x and has uniformly bounded diameter. Since $g\mathcal{F}_i^x, f(\mathcal{F}_i^x)$ coarsely coincide for $i = 1, 2$, we see that gx and $f(x)$ must coarsely coincide. Hence we get that the automorphism of $\mathcal{X}$ given by left-multiplication by g is uniformly close to the quasi-isometry f, as desired. □

6. FACTORED SPACES

Notation 6.1. Given $\mathfrak{U} \subseteq \mathfrak{S}$, let $\mathfrak{U}^{\sqsubseteq}$ be the collection of all $V \in \mathfrak{S}$ so that there exists $U \in \mathfrak{U}$ with $V \sqsubseteq U$. We let $\mathfrak{U} = \mathfrak{U}_{\mathcal{X}} \subset \mathfrak{S}$ denote the union of all cardinality–ν pairwise-orthogonal subsets of $\mathfrak{S}$. Let $\widehat{\mathcal{X}}$ be the factored space associated to $\mathfrak{U}^{\sqsubseteq}$, which is the space obtained from $\mathcal{X}$ by coning off all F_U for $U \in \mathfrak{U}^{\sqsubseteq}$ (as described in [BHS17a, Definition 2.1]). There exists a Lipschitz factor map $q = q_{\mathcal{X}} \colon \mathcal{X} \to \widehat{\mathcal{X}}$ by [BHS17a, Proposition 2.2].

$\widehat{\mathcal{X}}$ has a natural HHS structure with index set $\mathfrak{S} - \mathfrak{U}^{\sqsubseteq}$, by [BHS17a, Proposition 2.4].

Theorem 6.2. *Let $\mathcal{X}$ be an asymphoric HHS of rank ν. For any K, there exists Δ so that for all (K, K)–quasi-isometric embeddings $f \colon \mathbb{R}^\nu \to \mathcal{X}$, we have $\mathsf{diam}(q \circ f(\mathbb{R}^\nu)) \leqslant \Delta$.*

Proof. Observe that if $A \subset \mathcal{X}$, then $q(H_\theta(A))$ lies at uniformly bounded Hausdorff distance from $H_\theta(q(A))$ (where we take hulls in $\widehat{\mathcal{X}}$ in the second expression). In particular, if $\mathsf{diam}_{\widehat{\mathcal{X}}}(q(A)) \leqslant C$ for some C, then there exists $C' = C'(C, E, \theta)$ so that for any $B \subset H_\theta(A)$ we have $\mathsf{diam}_{\widehat{\mathcal{X}}}(q(B)) \leqslant C'$.

Hence, by Corollary 4.16, it suffices to prove that $\mathsf{diam}_{\widehat{\mathcal{X}}}\, q(\bigcup_{i=1}^N O_i) \leqslant C$, where the orthants O_i are as in the Corollary and $C = C(N, E, K, \mu_0)$. By the construction of q, it follows easily that there exists $C' = C'(\mu_0, E)$ such that $\mathsf{diam}_{\widehat{\mathcal{X}}}(q(O_i)) \leqslant C'$ for each i. By Proposition 6.6, it suffices to bound the diameter of $q(O_i \cup O_j)$ in the case where $O_i \tilde{\cap} O_j$ is a codimension-1 sub-orthant; this is done in Lemma 6.5. □

Before proceeding to the technical Lemmas and Propositions we needed to prove the above theorem, we state the following corollary which we consider the main result of this section.

Corollary 6.3. *Let $\mathcal{X},\mathcal{Y}$ be asymphoric HHS. Suppose that there exists D so that for each $U \in \mathfrak{U}_{\mathcal{X}}$ or $U \in \mathfrak{U}_{\mathcal{Y}}$, for any $x, y \in F_U$ there exists a bi-infinite (D,D)–quasi-geodesic containing x, y. Then for every quasi-isometry $f : \mathcal{X} \to \mathcal{Y}$ there exists a quasi-isometry $\hat{f} : \widehat{\mathcal{X}} \to \widehat{\mathcal{Y}}$ so that the diagram*

$$\begin{array}{ccc} \mathcal{X} & \xrightarrow{f} & \mathcal{Y} \\ {\scriptstyle q_{\mathcal{X}}}\downarrow & & \downarrow{\scriptstyle q_{\mathcal{Y}}} \\ \widehat{\mathcal{X}} & \xrightarrow[\hat{f}]{} & \widehat{\mathcal{Y}} \end{array}$$

commutes.

Proof. Since $\widehat{\mathcal{X}},\widehat{\mathcal{Y}}$ are just re-metrized copies of $\mathcal{X}$, $\mathcal{Y}$, we can take $\hat{f} = f$.

We now show that $\hat{f}$ is coarsely Lipschitz, and observe that the corresponding map for a quasi-inverse of f gives a coarsely Lipschitz inverse of $\hat{f}$.

By the definition of the metric on $\widehat{\mathcal{X}},\widehat{\mathcal{Y}}$ ([BHS17a, Definition 2.1]), we just have to verify that if x, y lie in some F_U for $U \in \mathfrak{U}^{\sqsubseteq}_{\mathcal{X}}$, then their images are uniformly close in $\widehat{\mathcal{Y}}$. By assumption, x, y lie close to a quasiflat with uniform constant, so that the conclusion follows from Theorem 6.2. □

Lemma 6.4. *There exists τ with the following property. Let O, O' be standard orthants in $\mathcal{X}$ with supports $\mathcal{U}_1,\mathcal{U}_2$. Suppose that $O\tilde{\cap}O'$ is a k–orthant whose support is $\mathcal{U}$. Then for each $x, y \in O \cup O'$ we have that any $U \in \mathfrak{S}$ with $\mathsf{d}_U(x,y) \geqslant \tau$ is either nested into some $U' \in \mathcal{U}_1 \cap \mathcal{U}_2$ or orthogonal to all $U' \in \mathcal{U}$.*

Proof. Recall that $O\tilde{\cap}O'$ coarsely coincides with $\mathfrak{g}_O(O')$ by Lemma 4.10 (and also with a standard orthant whose support is contained in $\mathcal{U}_1 \cap \mathcal{U}_2$, thereby describing $\mathcal{U}$).

The conclusion clearly holds if x, y both lie in either O or O' (by definition of standard product regions). We can then prove the lemma for $x \in \mathfrak{g}_O(O')$ and $y = \mathfrak{g}_{O'}(x)$, but in this case the conclusion follows from Lemma 1.19.(5). □

Lemma 6.5. *There exists $C = C(E,\mu_0)$ so that the following holds. Let O, O' be standard orthants with $O\tilde{\cap}O'$ a codimension-1 sub-orthant. Then $\mathsf{diam}_{\widehat{\mathcal{X}}}(q(O \cup O')) \leqslant C$.*

Proof. Let $x \in O, y \in O'$. Let $\mathcal{M} = \{U \in \mathfrak{S} : \mathsf{d}_U(x,y) \geqslant \tau\}$. By Lemma 6.4, each $U \in \mathcal{M}$ belongs to a set of pairwise-orthogonal elements of size ν (note that in the case that U is orthogonal to the intersection, this has maximal rank because of the fact that we are assuming the intersection has co-dimension-1). Hence $\mathsf{d}_U(q(x), q(y)) \leqslant \tau$ for all $U \in \mathfrak{S} - \mathfrak{U}$, so $q(x)$ is uniformly close to $q(y)$ by the uniqueness axiom. □

Proposition 6.6. *Suppose that the quasiflat $\mathcal{F}$ lies within finite Hausdorff distance of $\bigcup_{i=1}^{m} O_i$, where the O_i are standard orthants with $\mathsf{d}_{haus}(O_i, O_j) = \infty$ for $i \neq j$. Then for each pair of distinct orthants O_j, O_k there exists a sequence $j = j_0, \ldots, j_l = k$ so that the coarse intersection of O_{j_i} and $O_{j_{i+1}}$ is an $(\nu - 1)$–orthant.*

Proof. Passing to an asymptotic cone, we get a bilipschitz copy $\boldsymbol{\mathcal{F}}$ of $\mathbb{R}^\nu$ filled by bilipschitz copies $\boldsymbol{O}_i$ of $[0,\infty)^\nu$. The intersections of the $\boldsymbol{O}_i$ have some basic properties:

Lemma 6.7.

(1) The intersection of $\boldsymbol{O}_i$ and $\boldsymbol{O}_j$ is bilipschitz equivalent to $[0,\infty)^t$ for some $t = t(i,j)$.
(2) $t(i,j) = \nu - 1$ if and only if $\boldsymbol{O}_i$ and $\boldsymbol{O}_j$ coarsely intersect in an $(\nu - 1)$–orthant.

Proof. Recall that the coarse intersection of two standard orthants coarsely coincides with a standard k–orthant, as well as with the gate of one in the other (Lemma 4.10). We now show the following, which implies both statements: if the ultralimits $\boldsymbol{A}$, $\boldsymbol{B}$ of uniformly hierarchically quasiconvex sets have non-empty intersection, then their intersection is the ultralimit $\mathfrak{g}_{\boldsymbol{A}}(\boldsymbol{B})$ of the gates. By Lemma 1.19.(3), $\mathfrak{g}_{\boldsymbol{A}}(\boldsymbol{B})$ is contained in $\boldsymbol{A} \cap \boldsymbol{B}$ (this uses $\mathsf{d}(\boldsymbol{A}, \boldsymbol{B}) = 0$). Lemma 1.19.(6) implies that the other containment holds. □

Now, consider the subspace $X \subset \boldsymbol{\mathcal{F}}$ consisting of the union of all $\boldsymbol{O}_i \cap \boldsymbol{O}_j$ for i, j with $t(i, j) = \nu - 1$. Let $\mathcal{Y}$ be the set of all $\boldsymbol{O}_i \cap \boldsymbol{O}_j$ with $i \neq j$ and $t(i, j) < \nu - 1$. Let $Y = \bigcup_{O \in \mathcal{Y}} O$.

Lemma 6.8. *$\boldsymbol{\mathcal{F}} - Y$ is path-connected.*

Proof. In this proof, when referring to homology, we always mean homology with rational coefficients. The goal is to show $H_0(\boldsymbol{\mathcal{F}} - Y) = \mathbb{Q}$.

If $\dim \boldsymbol{\mathcal{F}} \leqslant 2$, then $\mathcal{Y}$ is a finite set (which is empty when $\dim \boldsymbol{\mathcal{F}} \leqslant 1$) and the claim is clear. Hence suppose that $\dim \boldsymbol{\mathcal{F}} \geqslant 3$. We argue by induction on $|\mathcal{Y}|$.

We first claim that for any $O \in \mathcal{Y}$ and any closed $O' \subset O$, $\boldsymbol{\mathcal{F}} - O'$ is path-connected and $H_1(\boldsymbol{\mathcal{F}} - O') = 0$. We use the fact that, for A, B closed homeomorphic subsets of $\mathbb{R}^\nu$ we have $H_*(\mathbb{R}^\nu - A) = H_*(\mathbb{R}^\nu - B)$, see e.g. [Dol93]. Hence, we can regard O as a coordinate orthant in $\mathbb{R}^\nu \cong \mathcal{F}$. Hence the claim holds for $O' = O$. The fact that $H_1(\boldsymbol{\mathcal{F}} - O') = 0$ follows from the fact that $H_1(\boldsymbol{\mathcal{F}} - O) = 0$, since a 1–cycle in $\boldsymbol{\mathcal{F}} - O'$ is homologous to one in $\boldsymbol{\mathcal{F}} - O$ by, for example, a transversality argument. The same holds for $H_0(\boldsymbol{\mathcal{F}} - O')$.

For the inductive step, let A be the union of all but one element of $\mathcal{Y}$, and let B be the remaining one. We have a Mayer-Vietoris sequence:

$$H_1(\boldsymbol{\mathcal{F}} - (A \cap B)) \to H_0(\boldsymbol{\mathcal{F}} - (A \cup B)) \to H_0(\boldsymbol{\mathcal{F}} - A) \oplus H_0(\boldsymbol{\mathcal{F}} - B) \to H_0(\boldsymbol{\mathcal{F}} - (A \cap B)) \to 0.$$

By the claim above, the first term is 0, the last term is $\mathbb{Q}$, and $H_0(\boldsymbol{\mathcal{F}} - B) = \mathbb{Q}$. By induction, $H_0(\boldsymbol{\mathcal{F}} - A) = \mathbb{Q}$. Hence $\boldsymbol{\mathcal{F}} - (A \cup B)$ is connected. □

We now finish the proof of Proposition 6.6.

Let $\boldsymbol{O}_j, \boldsymbol{O}_k$ be orthants. We will now produce a sequence $\boldsymbol{O}_j = \boldsymbol{O}_{j_0}, \ldots, \boldsymbol{O}_{j_l} = \boldsymbol{O}_k$ of orthants so that $t(j_i, j_{i+1}) = \nu - 1$ for $0 \leqslant i \leqslant l - 1$. Choose $\boldsymbol{x} \in \operatorname{Int}(\boldsymbol{O}_i), \boldsymbol{y} \in \operatorname{Int}(\boldsymbol{O}_j)$ and let $\sigma\colon [0, 1] \to \boldsymbol{\mathcal{F}} - Y$ be a path joining them, which is provided by Lemma 6.8. Let t_0 be the maximal t so that $\sigma(t) \in \boldsymbol{O}_j$. If $t_0 = 1$, then we take $l = 0$. Otherwise, there exists $\boldsymbol{O}_{j_1} \neq \boldsymbol{O}_j$ so that $\boldsymbol{O}_j \cap \boldsymbol{O}_{j_1}$ has dimension $\nu - 1$ and contains $\sigma(t_0)$. Now apply the same argument to $\sigma|_{[t_0,1]}$ and induct.

The sequence in the cone yields a sequence of orthants in the space with the desired property. □

References

[BF06] J. Brock and B. Farb. Curvature and rank of Teichmüller space. *American Jour. Math.*, 128(1):1–22, 2006.

[BHS15] Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: combination theorems and the distance formula. *arXiv:1509.00632*, 2015.

[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] J. Behrstock, M.F. Hagen, and A. Sisto. Hierarchically hyperbolic spaces I: curve complexes for cubical groups. *Geom. Topol.*, 21:1731–1804, 2017.

[BJN10] J. Behrstock, T. Januszkiewicz, and W. Neumann. Quasi-isometric classification of high dimensional right angled artin groups. *Groups, Geometry, and Dynamics*, 4(4):681–692, 2010.

[BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. *Geom. Topol.*, 16(2):781–888, 2012.

[BKS08] M. Bestvina, B. Kleiner, and M. Sageev. The asymptotic geometry of right-angled Artin groups. I. *Geometry & Topology*, 12(3):1653–1699, 2008.

[BKS16] Mladen Bestvina, Bruce Kleiner, and Michah Sageev. Quasiflats in CAT(0) 2-complexes. *Algebr. Geom. Topol.*, 16(5):2663–2676, 2016.

[BM08] J. Brock and H. Masur. Coarse and synthetic Weil–Petersson geometry: quasi-flats, geodesics, and relative hyperbolicity. *Geometry & Topology*, 12:2453–2495, 2008.

[BN08] J. Behrstock and W. Neumann. Quasi-isometric classification of graph manifold groups. *Duke Math. J.*, 141(2):217–240, 2008.

[Bow13] Brian Bowditch. Coarse median spaces and groups. *Pacific Journal of Mathematics*, 261(1):53–93, 2013.

[Bow14] Brian Bowditch. Uniform hyperbolicity of the curve graphs. *Pacific Journal of Mathematics*, 269(2):269–280, 2014.

[Bow15] Brian H Bowditch. Large-scale rigidity properties of the mapping class groups. *preprint*, 2015.

[Bro02] Jeffrey F. Brock. Pants decompositions and the Weil-Petersson metric. In *Complex manifolds and hyperbolic geometry (Guanajuato, 2001)*, volume 311 of *Contemp. Math.*, pages 27–40. Amer. Math. Soc., Providence, RI, 2002.

[Che00] Victor Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.

[CN05] Indira Chatterji and Graham Niblo. From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.*, 15(5-6):875–885, 2005.

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

[Dol93] Albrecht Dold. A simple proof of the Jordan-Alexander complement theorem. *Amer. Math. Monthly*, 100(9):856–857, 1993.

[Dru02] Cornelia Druţu. Quasi-isometry invariants and asymptotic cones. *International Journal of Algebra and Computation*, 12(01n02):99–135, 2002.

[EF97] A. Eskin and B. Farb. Quasi-flats and rigidity in higher rank symmetric spaces. *Jour. AMS*, 10(3):653–692, 1997.

[EMR] Alex Eskin, Howard Masur, and Kasra Rafi. Large scale rank of Teichmuller space. *Duke Math. J.* To appear.

[FLS15] Roberto Frigerio, Jean-François Lafont, and Alessandro Sisto. Rigidity of high dimensional graph manifolds. *Astérisque*, (372):xxi+177, 2015.

[Gro87] M. Gromov. Hyperbolic groups. In S. Gersten, editor, *Essays in group theory*, volume 8 of *MSRI Publications*. Springer, 1987.

[Ham07] U Hamenstädt. Geometry of the mapping class group iii: Quasi-isometric rigidity. preprint. *arXiv preprint math.GT/051242*, 2007.

[HP98] Frédéric Haglund and Frédéric Paulin. Simplicité de groupes d'automorphismes d'espaces à courbure négative. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 181–248. Geom. Topol. Publ., Coventry, 1998.

[HS] Mark F Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. arXiv:1609.01313.

[Hua14a] Jingyin Huang. Quasi-isometry rigidity of right-angled Artin groups I: the finite out case. *to appear in Geometry & Topology, arXiv:1410.8512*, 2014.

[Hua14b] Jingyin Huang. Top dimensional quasiflats in $CAT(0)$ cube complexes. *to appear in Geometry & Topology, arXiv:1410.8195*, 2014.

[Hua16] Jingyin Huang. Quasi-isometry classification of right-angled Artin groups II: several infinite out cases. *arXiv:1603.02372*, 2016.

[HW14] G. C. Hruska and Daniel T. Wise. Finiteness properties of cubulated groups. *Compos. Math.*, 150(3):453–506, 2014.

[Iva97] Nikolai V Ivanov. Automorphisms of complexes of curves and of teichmuller spaces. *International Mathematics Research Notices*, 1997(14):651–666, 1997.

[KL97a] M. Kapovich and B. Leeb. Quasi-isometries preserve the geometric decomposition of Haken manifolds. *Invent. Math.*, 128(2):393–416, 1997.

[KL97b] B. Kleiner and B. Leeb. Rigidity of quasi-isometries for symmetric spaces and euclidean buildings. *IHES Publ. Math.*, 86:115–197, 1997.

[Mor24] H. M. Morse. A fundamental class of geodesics on any closed surface of genus greater than one. *Trans. AMS*, 26(1):25–60, 1924.

[Mos73] G. D. Mostow. *Strong Rigidity of Locally Symmetric Spaces*. Number 78 in Annals of Math. Studies. Princeton Univ. Press, 1973.

[MS13] Howard Masur and Saul Schleimer. The geometry of the disk complex. *J. Amer. Math. Soc.*, 26(1):1–62, 2013.

[Nic04] Bogdan Nica. Cubulating spaces with walls. *Algebr. Geom. Topol.*, 4:297–309, 2004.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[Zei16] Rudolf Zeidler. Coarse median structures and homomorphisms from Kazhdan groups. *Geometriae Dedicata*, 180(1):49–68, 2016.

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

Dept. of Pure Maths and Math. Stat., University of Cambridge, Cambridge, UK
E-mail address: markfhagen@gmail.com

ETH, Zürich, Switzerland
E-mail address: sisto@math.ethz.ch