

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 in a standard product region. For hierarchically hyperbolic groups, this coincides with the maximal dimension of a quasiflat. 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; in the context of certain $\text{CAT}(0)$ cubical groups, our result handles novel special cases, including right-angled Coxeter groups.

An important ingredient in the proof, which we expect will have other applications, is that the *hull* of any finite set in an HHS is quasi-isometric to a $\text{CAT}(0)$ cube complex of dimension bounded by the rank (if the HHS is a $\text{CAT}(0)$ 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 HHSs induces a quasi-isometry between certain *factored spaces*, which are simpler HHSs. 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, which, once we’ve established our general quasiflats theorem, uses simpler combinatorial arguments than in previous proofs.

CONTENTS

Introduction	2
1. Background	11
2. Cubulation of hulls	25
3. Quasiflats and asymptotic cones	36
4. Orthants and quasiflats	41
5. Induced maps on hinges: mapping class group rigidity	48
6. Factored spaces	53
References	56

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 coarsely 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 a quasigeodesic in a hyperbolic space, a quasiflat need not lie close to any one flat. However, generalizing Mostow’s result, Eskin–Farb and Kleiner–Leeb each proved that, 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 explain the structure of quasiflats in a broad class of spaces and groups with a property called *hierarchical hyperbolicity* [BHS17b, BHS19, BHS17a]. Hierarchical hyperbolicity captures the coarse nonpositive curvature visible in many important groups and spaces, including mapping class groups, right-angled Artin groups, many CAT(0) cube complexes, most 3-manifold groups, Teichmüller space (in any of the standard metrics), etc.

Hierarchical hyperbolicity generalizes, and was inspired by, theorems about the mapping class group established by Masur–Minsky [MM00], Behrstock [Beh06], and others. Motivation also comes from Kim–Koberda’s work towards obtaining an analogue of some of those mapping class group results in the setting of right angled Artin groups [KK14]. To approach other problems, some features of the mapping class group were axiomatized by Bestvina–Bromberg–Fujiwara to great effect [BBF15, BBF19].

The class of hierarchically hyperbolic spaces is preserved by quasi-isometries, and also includes many examples not on the preceding list: one can readily produce 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 [BHS19, BR17]. Groups that are hyperbolic relative to hierarchically hyperbolic groups are again hierarchically hyperbolic [BHS19]. It is shown in [BHS17a] that suitable small-cancellation quotients of hierarchically hyperbolic groups are again hierarchically hyperbolic.

This article establishes a relationship between some of these examples: in particular, we show that these spaces all admit a very strong local approximation by CAT(0) cube complexes (Theorem F). This allows us to use cubical techniques in new settings. For example, it enables application of cubical geometry to mapping class groups.

Even for CAT(0) cube complexes our approximation provides new information. The reason is that Theorem F allows one to approximate convex hulls of finite sets in an HHS by finite CAT(0) cube complexes, and if the ambient HHS is a CAT(0) cube complex, the dimension of the approximating complex — which is bounded by the rank — can be much lower than the dimension of the ambient complex. This is essential for our applications to quasiflats.

Our techniques are intrinsic to the category of hierarchically hyperbolic spaces, in the sense that the arguments in this paper couldn’t be carried out strictly in the context of any of the motivating examples alone, for example CAT(0) cube complexes or mapping class groups.

Formal definitions and relevant properties of hierarchically hyperbolic spaces (HHSs) will be given below in Section 1. For now, we recall that a *hierarchically hyperbolic space* consists of: a quasigeodesic metric space, \mathcal{X} ; an *index set*, \mathfrak{S} ; a hyperbolic space $\mathcal{C}U$ for each $U \in \mathfrak{S}$; some relations between elements of the index set and maps between the associated hyperbolic spaces. There are also projections $\mathcal{X} \rightarrow \mathcal{C}U, U \in \mathfrak{S}$, and various axioms governing all of this data.

Before stating the main theorem, we informally recall a few geometric features of HHSs:

- Any HHS \mathcal{X} contains certain *standard product regions*, in which each of the (boundedly many) factors is itself an HHS.

In mapping class groups, these are products of mapping class groups of pairwise disjoint subsurfaces (for example the subgroup generated by Dehn twists along disjoint annuli). In CAT(0) cube complexes, these are certain convex subcomplexes that split as products (for example in the Salvetti complex of a right-angled Artin group, they are subcomplexes associated to subgraphs of the defining graph that decompose as joins).

- 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 $[0, \infty) \rightarrow \mathcal{X}$ is a quasi-isometric embedding $[0, \infty)^k \rightarrow \mathcal{X}$ 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* $U_1, \dots, U_\nu \in \mathfrak{S}$ for which each CU_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 HHSs, including all hierarchically hyperbolic groups. Under this condition, Theorem 1.15 implies that the rank is a quasi-isometry invariant.

Quasiflats. Understanding the structure of quasiflats in a given metric space or group is often critical in understanding the geometry of that space.

An early version of a “quasiflats theorem” is Mostow’s result that in a rank-one symmetric space, any quasi-geodesic lies within a uniformly bounded distance of a geodesic [Mos73].

A well known generalization of this is due to Schwartz, who proved that the image of any quasi-isometric embedding of $\mathbb{R}^n, n \geq 2$ into a non-uniform lattice in a rank-one symmetric space lies within a uniformly bounded distance of a horosphere [Sch95]. (Actually, Schwartz proved a more general result, namely that the image of any quasi-isometric embedding from a space whose asymptotic cone doesn’t have cut-points into a “neutered space” lies uniformly close to a horosphere; he credits unpublished work of Gersten for the case of Euclidean space.)

Schwartz’s result was generalized by Druţu–Sapir, who replaced the target space by an arbitrary relatively hyperbolic space and showed that the image of the quasi-isometric embedding lies uniformly close to a peripheral subspace [DS05].

This result was in turn generalized by Behrstock–Druţu–Mosher, who weakened the hypothesis on the domain to allow any space which is itself not relatively hyperbolic [BDM09].

As noted above, Eskin–Farb and Kleiner–Leeb proved that, in a higher-rank symmetric space, an arbitrary quasiflat must lie within finite distance of a finite number of flats [EF97, KL97b].

As discussed further below, there has been much work toward quasiflats theorems in other contexts. We now state our main result in this direction, explain some consequences, and describe interactions with related work. At the end of the introduction we sketch the proof.

Theorem A (Quasiflats Theorem for HHSs). *Let \mathcal{X} be an HHG of rank ν . Let $f: \mathbb{R}^\nu \rightarrow \mathcal{X}$ be a quasi-isometric embedding. Then there exist standard orthants $Q_i \subseteq \mathcal{X}, i = 1, \dots, k$, so that $d_{\text{haus}}(f(\mathbb{R}^\nu), \cup_{i=1}^k Q_i) < \infty$. 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 and discuss related results.

Mapping class groups have been studied for the past two decades using tools introduced by Masur–Minsky [MM00]; these tools were then further developed in [Beh06, BKMM12] and elsewhere. The results of [Beh06, BKMM12, MM00] together show that mapping class groups are hierarchically hyperbolic; see [BHS19, Theorem 11.1] for details. Theorem A, applied to mapping class groups, 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 lies within a uniformly bounded distance from a finite union of standard flats.*

Although the structure of quasiflats in the mapping class group was unsettled, numerous prior results were obtained in pursuit of the resolution of this conjecture. One partial result in this direction was Behrstock–Minsky’s theorem that \mathbb{R}^n can only be quasi-isometrically embedded in a given mapping class group if n is at most the complexity of the surface [BM08a]. This established the dimension of the top-dimensional quasiflats in the mapping class group.

Also significant are a number of results which give some local control of top-dimensional quasiflats in the mapping class group. In particular, see results of Behrstock–Kleiner–Minsky–Mosher [BKMM12], Bowditch [Bow18b], and Eskin–Masur–Rafi [EMR17]. Although those prior results yield some control over quasiflats, Theorem A is the first to completely describe the structure of quasiflats in the mapping class group. As we will describe in more detail later, we use some of the tools developed by Bowditch in [Bow18b] in our proof of Theorem A.

Outside of the setting of groups, we apply Theorem A to the Weil-Petersson metric on Teichmüller space, which is an asymphoric HHS by virtue of Brock’s theorem that the pants graph is quasi-isometric to the Weil-Petersson metric [Bro03] and results of [Beh06, BKMM12, MM00]; for details see [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, answering his question in the affirmative.

Corollary C (Affirmation of Brock’s Question: Quasiflats theorem for Weil-Petersson metric). *Any top-dimensional quasiflat in the Weil-Petersson metric on Teichmüller space lies within a uniformly bounded distance from a finite union of standard flats.*

The previously answered cases of Brock’s question were: in the rank one cases, a positive answer comes from Brock–Farb’s result that the space is hyperbolic [BF06]; in the three rank-two cases, Brock–Masur proved that the space is relatively hyperbolic and thus that each quasiflat is contained in a single peripheral subset [BM08b, Theorem 3]. In the general case, there were partial results providing coarse local control; in particular, there are theorems about flats being locally contained in linear size neighborhoods of standard flats, e.g., [BKMM12, Theorem 8.5] and [EMR17, Theorem A].

Fundamental groups of non-geometric 3-manifolds are HHSs of rank 2 [BHS19, Theorem 10.1]. For these groups, Theorem A allows us to recover the following quasiflats theorem, which was first established by Kapovich–Leeb:

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

Some quasiflat theorems have previously been obtained for CAT(0) spaces satisfying particular conditions.

One such result is due to Bestvina–Kleiner–Sageev, who proved that, for two-dimensional, proper, piecewise Euclidean CAT(0) complexes admitting cocompact group actions, every

two-dimensional quasiflat lies within finite distance from a subset which is locally flat outside a compact set [BKS16].

Generalizing that result, Huang proved that in an N -dimensional CAT(0) cube complex, every N -dimensional quasiflat lies within finite distance from a finite union of standard orthants [Hua14b]. The following corollary of Theorem A generalizes the above noted theorems of [BKS16] and [Hua14b] in certain cocompact cases:

Corollary E (Quasiflats theorem for cubulated groups with factor systems). *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 $f: \mathbb{R}^\nu \rightarrow \mathcal{X}$ be a quasi-isometric embedding. Then $d_{\text{haus}}(f(\mathbb{R}^\nu), \cup_{i=1}^k Q_i) < \infty$, where each Q_i can be chosen to be:*

- an ℓ^1 -isometrically embedded copy of the standard cubical tiling of $[0, \infty)^\nu$, or
- if \mathcal{X} admits a cocompact group action, 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 (in the sense of Haglund-Wise [HW08]) group actions admit factor systems. More generally, it is shown in [HS18] 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 (indeed, we have conjectured that all cubical groups admit factor systems, see [BHS19, Conjecture A]).

Proof of Corollary E. 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 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 $o: \prod_{i=1}^\nu \gamma_i \rightarrow \mathcal{X}$ with γ_i a combinatorial geodesic ray, there is a CAT(0) orthant o' with $d_{\text{haus}}(\text{im}(o), o') \leq N$. For each i , let \mathcal{Y}_i be the convex hull of γ_i , i.e., the intersection of all combinatorial half-spaces containing γ_i . Then the hull of $\text{im}(o)$ 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 $[0, m]^2 \rightarrow \mathcal{Y}_i$, and thus an ℓ_1 -isometric embedding $[0, m]^2 \times [0, \infty)^{\nu-1} \rightarrow \mathcal{X}$. Cocompactness would then allow us to produce a $(\nu+1)$ -dimensional cubical orthant in \mathcal{X} , which is impossible by our choice of ν . \square

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 exist hyperbolic (and hence rank one) cubulated groups, whose associated CAT(0) cube complexes have arbitrarily large dimension.

In this sense, this corollary is stronger than the main result in [Hua14b]; our result applies even if the dimension is larger than the rank. On the other hand, in practice, the construction of a factor system relies on a geometric group action, a hypothesis not needed in the context of [Hua14b].

Although our results, applied in the cubical case, generalize some of those of [Hua14b], our proof is obtained by passing from the hierarchically hyperbolic space setting to a CAT(0) cube complex where the dimension equals the rank and then using Huang's theorem. Specifically,

we use the *cubical approximations*, discussed immediately below, to construct a CAT(0) cube complex to which we can apply Huang’s result, [Hua14b, Theorem 1.1]. So, Huang’s result is a crucial ingredient in our work.

Approximating with cube complexes. A key insight in the geometry of hyperbolic spaces is that in certain respects, they “coarsely look like trees”; Gromov, in his famous treatise on the subject, introduced a number of ways in which this idea can be made precise [Gro87a]. One such statement is: in a hyperbolic space, the coarse convex hull of any finite set of points can be uniformly approximated by a geodesic tree [Gro87a, §6.2 Geodesic trees].

It is now well understood that CAT(0) cube complexes are a natural generalization of trees. Two important aspects of this idea are:

- In a simplicial tree, the midpoint of any edge separates the tree into two complementary components. In a CAT(0) cube complex, the midpoint of each edge is contained in a *hyperplane*, a codimension–1 subspace with exactly two complementary components. The revolutionary work of Sageev [Sag95], elaborated later in [CN05, Nic04, HW14], shows that very general set-theoretic data — a *wallspace*, i.e. a set equipped with a suitable collection of bipartitions — determines a CAT(0) cube complex in a canonical way. We need this in Section 2.
- In a simplicial tree, any three vertices determine a unique geodesic tripod consisting of three geodesics, each of which joins two of the given points. The intersection of the three geodesics is a single vertex, the *median* of the three points. Generalizing this, one obtains the class of *median graphs*, i.e. graphs where each triple of vertices spans at least one metric tripod, all of which have the same center. Chepoi showed [Che00] that there is a bijective correspondence between one-skeleta of CAT(0) cube complexes and median graphs. The median viewpoint on CAT(0) cube complexes is very useful, and we adopt it in various ways in this paper.

In Section 2, we generalize Gromov’s theorem about hyperbolic groups to the setting of hierarchically hyperbolic spaces. Roughly, the theorem we prove establishes that the “convex hull” of a finite set A , denoted $H_\theta(A)$, is approximated by a finite CAT(0) cube complex.

This result provides a new tool for studying hierarchically hyperbolic spaces. Indeed, it is one of the key innovations which allows us to apply Huang’s theorem about quasi-flats in CAT(0) cube complexes [Hua14b] to prove Theorem A about quasiflats in HHSs. Further, we expect that Theorem F will have a number of applications beyond those of this paper. A sketch of the proof of this result is provided later in the introduction.

Theorem F (Approximation of convex hulls in HHSs 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 –*quasimedial* (C, C) –*quasi-isometry* $\mathfrak{p}_A: \mathcal{Y} \rightarrow H_\theta(A)$.*

A new proof of the preceding theorem, in a slightly more general context, was given by Bowditch [Bow18a], motivated by an early version of this paper.

Any HHS is coarse median in the sense of Bowditch [Bow13], as shown in [BHS19, Section 7]. In the coarse median setting, there are several interesting precursors to our theorem. One which was particularly inspirational to us was Bowditch’s result that in the asymptotic cone of a finite rank coarse median metric space, any top-dimensional closed Euclidean flat is cubulated, see [Bow18b, Proposition 1.2]. We will use Bowditch’s result about cubulating top-rank Euclidean subsets in a complete median space in order to apply our result about cubulating arbitrary finite sets in an HHS.

In the coarse median setting, one has, by definition, “cubical approximations of finite sets,” and there is also a nice metric approximation result given by Zeidler [Zei16, Theorem

6.2]. The approximation given by Theorem F has stronger properties, as it provides an approximation of the entire convex hull, whereas the *quasimediant 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 $n \geq 0$. 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 [BHS19]; see Corollary 2.15. This also leads to a characterization of hierarchically hyperbolic spaces which are hyperbolic, which we establish as Corollary 2.16.

Induced quasi-isometries on factored spaces and quasi-isometric classification. In [BHS17a, §2], 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. Factored spaces are HHSs themselves, with respect to a substructure of the original HHS structure. Factored spaces are central in the proof of finite asymptotic dimension [BHS17a].

A notable naturally-occurring example is that the Weil-Petersson metric on Teichmüller space is quasi-isometric to a factored space of the corresponding mapping class group. In any HHS, we proved in [BHS17a, Corollary 2.9] that there exists a factor space which is quasi-isometric to \mathcal{CS} for the \sqsubseteq -maximal element S (e.g., for the mapping class group of a surface S then \mathcal{CS} is the curve graph of 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{CS} 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 $f: \mathbb{R}^\nu \rightarrow \mathcal{X}$ satisfies $\text{diam}_{\mathcal{CS}}(\pi_S(f(\mathbb{R}^\nu))) \leq L$.*

In Corollary 6.3, we prove that any quasi-isometry between HHSs (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 1.15). 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, BJS10, BKS08].

Induced automorphisms of combinatorial data and quasi-isometric rigidity. The Quasiflats Theorem provides a powerful tool for proving quasi-isometric rigidity results for various HHSs, e.g. 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.

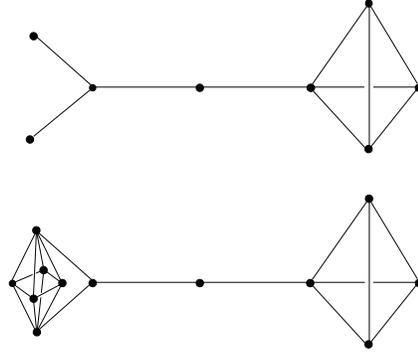


FIGURE 1. The right-angled Artin groups associated to the two graphs both have rank 4. However, the 4-dimensional flats get collapsed in the corresponding factored spaces, leaving only 2-dimensional flats in the case of the first RAAG, while there are 3-dimensional flats that persist in the case of the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

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 5.7 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 5.7 to give a new proof of quasi-isometric rigidity of mapping class groups which, once we have established Theorem A, requires simpler combinatorial considerations than previous proofs, cf. [BKMM12, Bow18b, Ham07].

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

Theorem 5.7 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 5.7 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 $star(\emptyset) = \Gamma$ (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 $\Lambda \subseteq \text{link}(\Lambda')$. 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 $\text{Aut}(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$ of $(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$. When is every element of $\text{Aut}(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$ induced by left multiplication by an element of B_Γ ? When is every element of $\text{Aut}(\mathfrak{S}_\Gamma, \sqsubseteq, \perp)$ “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 5.7 states that, under three natural assumptions, a quasi-isometry $f: (\mathcal{X}, \mathfrak{S}) \rightarrow (\mathcal{Y}, \mathfrak{T})$ 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 HHSs 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.

Sketch of the proof of the Cubulation of Hulls Theorem. We provide here a sketch of the proof of Theorem F (Approximation of convex hulls in HHSs by CAT(0) cube complexes). This is one of the main tools we develop in this paper, allowing one to use ideas from the world of cube complexes to study HHSs. This result plays a crucial role in the proof of Theorem A (Quasiflats Theorem for HHSs). The full proof is given in Section 2.

A hierarchically hyperbolic space can be roughly thought of as a subset of the product of a (typically infinite) collection of hyperbolic spaces. This subset has the property that its projection to any direct product of two factors is surjective if and only if those two factors are “orthogonal.” This allows one to move back and forth between properties of the HHS, \mathcal{X} , and properties in the associated hyperbolic spaces, $\{\mathcal{C}U\}$. Here is a construction from [BHS19, §6] that exploits this point of view.

Given a set of points in \mathcal{X} , one can build the “hull” of that set by looking at the projections of that finite set of points to each of the associated hyperbolic spaces, $\mathcal{C}U$, taking coarse convex hulls in the $\mathcal{C}U$, and then looking at the points in \mathcal{X} that in each $\mathcal{C}U$ project close to the hull.

The *realization theorem* [BHS19, Theorem 3.1] gives, roughly, a characterization of points in the product of the hyperbolic spaces that lie in the image of \mathcal{X} , in terms of *consistency conditions* (in the mapping class group context, one such condition is given by [Beh06, Theorem 4.3]). In this paper we rely on the construction of hulls and the realization theorem in an essential way.

Following Sageev, the main method to cubulate a space is to explicitly build *walls*, that is, “codimension-one” subspaces which separate the space. For hulls of a finite set of points in a hierarchically hyperbolic space, this can be done in the following manner. Consider a finite set of points and their projections to each hyperbolic space $\mathcal{C}U$. By Gromov’s theorem, in a hyperbolic space the convex hull of a finite set of points can be uniformly approximated by a geodesic tree [Gro87a, §6.2 Geodesic trees].

Taking an appropriately dense collection of points in each such geodesic tree and considering their inverse images in \mathcal{X} , one obtains walls that can be used to construct a CAT(0) cube complex. One needs to verify that this actually works as needed to prove Theorem F. In particular, a key point is to show that any vertex in the CAT(0) cube complex “corresponds” to a point in the hull in \mathcal{X} of the finitely many points. Establishing this requires a careful analysis of the “consistency conditions”, with the aim of invoking the aforementioned realization theorem, [BHS19, Theorem 3.1].

Sketch of the proof of the Quasiflats Theorem. An important ingredient in our proof of Theorem A (Quasiflats Theorem for HHSs) is Huang’s result [Hua14b, Theorem 1.1] which classifies n -dimensional quasiflats in n -dimensional CAT(0) cube complexes (for emphasis: in Huang’s theorem the dimension of the quasiflats under consideration coincides with that of the CAT(0) cube complex).

We prove Theorem A by constructing an appropriate CAT(0) cube complex to which we can apply Huang’s theorem.

The first step is to use a result of Bowditch [Bow18b, Proposition 1.2] about “local cubulations” of top-dimensional flats in median metric spaces. The median spaces in question are (bilipschitz equivalent to) asymptotic cones of the HHS \mathcal{X} . This allows us to show that any finite ball in a quasiflat is coarsely contained in the hull of a *uniformly* bounded number of points.

By our cubulation of hulls theorem discussed above, Theorem F, we know that any such hull is uniformly quasi-isometric to a CAT(0) cube complex. Taking an ultralimit of these CAT(0) cube complexes, we obtain a finite dimensional CAT(0) cube complex which quasi-isometrically embeds in our HHS, and the quasiflat is contained in a bounded neighborhood of the image of the quasi-isometric embedding.

By construction and Theorem F, the CAT(0) cube complex we build has the same dimension as the quasiflat, thus allowing us to apply Huang’s result [Hua14b, Theorem 1.1]. This finishes the proof since one can show that the orthants in the CAT(0) cube complex we construct correspond to standard orthants in the original HHS.

As can be seen in this sketch, our theorem relies on Huang’s result [Hua14b, Theorem 1.1] in an essential way. We also note that because the CAT(0) cube complex we construct is built using Theorem F, this complex always has the same dimension as its top-dimensional quasiflats. So, although our argument factors through Huang’s theorem, our result extends Huang’s in the setting of cocompact CAT(0) cube complexes with factor systems. For instance, our theorem applies to arbitrary right-angled Coxeter groups, even though the dimension of the CAT(0) cube complex associated to a right-angled Coxeter is typically (much) larger than the dimension of the quasiflats it contains.

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 3.9, 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 HHSs, and then prove Theorem A. We also prove 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 5.7. 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. We first prove Theorem 6.2 and then deduce Corollary 6.3, which is about induced quasi-isometries of factored spaces.

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 $\sigma\upsilon\mu\phi\omicron\rho\acute{\alpha}$ and roughly meaning “mishap averted.”

We thank Ela Behrstock for drawing two of the figures (Figure 2 and Figure 4). We 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. We thank Brian Bowditch, Harry Petyt, and Jacob Russell for several helpful comments. We

also thank the numerous referees for very useful comments; in particular, the 87 enumerated comments of one of the referees helped us to significantly improve the exposition.

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.

Behrstock was supported by NSF grant DMS-1710890. Hagen was supported by EPSRC grant EP/L026481/1.

Sisto was partially supported by the Swiss National Science Foundation (grant #182186).

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.

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 [BHS19]. 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 $\pi_U : \mathcal{X} \rightarrow \mathcal{C}U$;
- 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 $U \pitchfork V$ or $U \pitchfork V$, a uniformly bounded set ρ_V^U ;
- for $U \sqsubset V$, a coarse map $\rho_V^U : \mathcal{C}V \rightarrow \mathcal{C}U$.

Definition 1.1 of [BHS19] consists of several axioms governing this data. The definition and basic properties of HHSs were first laid out in [BHS17b]; below we list [BHS19] as the primary reference since a few of the properties were first established there and this provides for unified notation. The properties of HHSs which are central to this article are listed below.

Remark (QI invariance). As explained in [BHS19, Proposition 1.10], the property of being a hierarchically hyperbolic space is preserved under quasi-isometries. If $(\mathcal{X}, \mathfrak{S})$ is a hierarchically hyperbolic space and $f : \mathcal{X}' \rightarrow \mathcal{X}$ is a quasi-isometry, then $(\mathcal{X}', \mathfrak{S})$ is an HHS; where the structure in \mathcal{X} is obtained by replacing each projection $\pi_U, U \in \mathfrak{S}$ by $\pi_U \circ f$.

The first property 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

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

- if $V \sqsubset W$, then

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

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

Remark (Consistent tuples). The consistency axiom has a sort of converse, the *realization theorem* stated below. The idea is that the projection maps $\pi_U, U \in \mathfrak{S}$ allow us to think of points in \mathcal{X} as tuples in $\prod_{U \in \mathfrak{S}} \mathcal{C}U$. The consistency axiom imposes conditions on which tuples can be in the image of the map $\mathcal{X} \rightarrow \prod_{U \in \mathfrak{S}} \mathcal{C}U$ given by the projections. The realization theorem says that the consistency conditions actually (coarsely) characterize tuples in the image of \mathcal{X} . The precise statement is Theorem 1.2, which is formulated using the notion of a *consistent tuple*, which we now define.

Fix a constant $\kappa \geq 0$. Given a tuple $(b_U)_{U \in \mathfrak{S}} \in \prod_{U \in \mathfrak{S}} \mathcal{C}U$, we say that $(b_U)_{U \in \mathfrak{S}}$ is κ -consistent if it satisfies the conditions from Axiom 1.1, except with each occurrence of $\pi_V(x)$ (resp. $\pi_W(x)$) replaced by b_V (resp. b_W). So, the axiom says that tuples in the image of \mathcal{X} are E -consistent.

Now we can state the realization theorem:

Theorem 1.2 (Realization of consistent tuples; [BHS19]). *For each $\kappa \geq 1$ there exist $\theta_e, \theta_u \geq 0$ such that the following holds. Let $\vec{b} \in \prod_{W \in \mathfrak{S}} 2^{\mathcal{C}W}$ be κ -consistent ([BHS19, Definition 1.17]); for each W , let b_W denote the $\mathcal{C}W$ -coordinate of \vec{b} .*

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

The realization theorem is one of what we see as three foundational theorems about HHSs. The other two are closely related: the *distance formula* and the existence of *hierarchy paths*.

The distance formula provides a way to 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 $\{\mathcal{C}U\}_{U \in \mathfrak{S}}$.

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

Theorem 1.3 (Distance Formula; [BHS19]). *Let (X, \mathfrak{S}) be hierarchically hyperbolic. Then there exists s_0 such that for all $s \geq s_0$ there exist constants K, C such that for all $x, y \in \mathcal{X}$,*

$$d_{\mathcal{X}}(x, y) \asymp_{K,C} \sum_{W \in \mathfrak{S}} \{\!\{d_W(x, y)\}\!\}_s.$$

Remark (Uniqueness axiom). 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 special case is the *uniqueness axiom*, which is part of the definition of a hierarchically hyperbolic space [BHS19, Definition 1.1.(9)]. There are various places in Section 2 where we apply the distance formula, but could probably get away with just using the uniqueness axiom. In fact, since we initially posted this paper, Bowditch has given an independent proof of Theorem F, not using the distance formula. One can then deduce the distance formula from Theorem F, which Bowditch does in [Bow18a].

Hierarchy paths are quasi-geodesics in the HHS whose projections to each associated hyperbolic space are (coarsely) monotone. Any two points can be joined by a hierarchy path:

Theorem 1.4 (Existence of Hierarchy Paths; [BHS19]). *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 between $\pi_U(x)$ and $\pi_U(y)$ in $\mathcal{C}U$ for each $U \in \mathfrak{S}$.*

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

assertion is the *bounded geodesic image* axiom for an HHS [BHS19, Definition 1.1.(6)] and the second assertion follows easily from the first, together with Axiom 1.1; for ease of reference we record this here as:

Lemma 1.5. (*Bounded geodesic image*) *Let \mathcal{X} be a hierarchically hyperbolic space. There exists B so that the following holds. Let $W \in \mathfrak{S}$, $V \sqsubset W$. Suppose that γ is a geodesic in CW with $\gamma \cap \mathcal{N}_B(\rho_W^V) = \emptyset$. Then $\text{diam}_{CV}(\rho_V^W(\gamma)) \leq B$.*

Moreover, suppose $x, y \in \mathcal{X}$ and that there exists a geodesic γ in CW from $\pi_W(x)$ to $\pi_W(y)$ so that $\gamma \cap \mathcal{N}_B(\rho_W^V) = \emptyset$. Then $d_V(x, y) \leq B$.

Another part of the definition of a hierarchically hyperbolic space is the *large links axiom* (Definition 1.1.(7) in [BHS19]). It says roughly that, if $x, y \in \mathcal{X}$ and $V \in \mathfrak{S}$, then the number of $U \sqsubset V$ on which x, y have very different projections, and U is \sqsubseteq -maximal with those properties, can be bounded in terms of $d_V(x, y)$. Typically, one does not apply the large links axiom directly. Instead, one uses a consequence, Lemma 2.5 of [BHS19], which we call “passing up large projections.” We will use a variant of that lemma, which we state presently (it is applied in an essential way in the proof of Lemma 2.10, which is part of the proof of Theorem F).

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 $C \geq 0$ 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 $d_{V_i}(x, y) \geq E$. Then there exists $W \in \mathfrak{S}_V$ and i, j so that $V_i, V_j \sqsubset W$ and $d_W(\rho_W^{V_i}, \rho_W^{V_j}) \geq C$.*

Example 1.7. Since the statement of the preceding lemma is somewhat opaque, we now give an example before proceeding to the proof. Let \mathcal{X} be the Cayley graph of the free group on generators a, b . We can make \mathcal{X} an HHS by taking \mathfrak{S} to consist of all left cosets of all subgroups generated by subsets of $\{a, b\}$. The space $\mathcal{C}\langle a \rangle$ is just \mathbb{R} , and similarly for $\mathcal{C}\langle b \rangle$. The space $\mathcal{C}\langle a, b \rangle$ is obtained from \mathcal{X} by coning off each coset in \mathfrak{S} .

Consider the path $w = (a^E b^E)^N$, for some $E \geq 1$. Then there are N cosets of $\langle a \rangle$ and $\langle b \rangle$ on which the endpoints of the above path have projections lying at distance E . For any C , by making N sufficiently large, we see that the coset $\langle a \rangle$ and the coset $w b^{-E} \langle b \rangle$ are at least C -distant in $\mathcal{C}\langle a, b \rangle$ and hence satisfy the conclusion of the lemma.

Proof of Lemma 1.6. First of all, we choose constants. Let $B \geq 1$ 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 $U \sqsubset V$. Moreover, suppose $B \geq D$, for D as in Theorem 1.4, 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 CU .

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 $\leq k - 1$. 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 $|\{V_i\}| \geq N(k + 1)$ (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 $\max_{i,j} \{d_V(\rho_V^{V_i}, \rho_V^{V_j})\} \geq C$, 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

- $d_V(x', y') \leq C + 100B$,

- $\pi_V(x'), \pi_V(y')$ 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, $d_{V_i}(x', y') \geq 100kB$, since $d_{V_i}(x', y')$ is approximately equal to $d_{V_i}(x, y)$.

The large link axiom ([BHS19, Definition 1.1.(6)]) implies that there exists $K = K(C + 100B)$ and T_1, \dots, 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 $N(k + 1) \geq KN(k)$, there exists j so that $\geq N(k)$ elements of $\{V_i\}$ are nested into T_j . By the induction hypothesis, we are done. \square

1.1.1. *A few more basic HHS notions.* We now collect a few more basic notions about HHSs that will be used throughout the paper.

First, each of the HHS axioms (and their variants stated above) involves some constants, which are taken to be part of the HHS structure $(\mathcal{X}, \mathfrak{S})$. For the sake of sanity, where possible, we can assume these constants are all the same:

Notation 1.8 (Naming constants). In the remainder of the paper, following [BHS19, Remark 1.6], we fix a constant E larger than each of the constants in [BHS19, Definition 1.1] and also satisfying the conclusion of Lemma 1.6, Lemma 1.5, and Axiom 1.1.

Given $x, y \in \mathcal{X}$, it is convenient to consider the subset of \mathfrak{S} on whose associated hyperbolic spaces x, y project far apart, where ‘‘far’’ is determined by some threshold, generally specified in advance independently of x, y :

Definition 1.9 (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 $d_U(x, y) > \theta$. Denote by $\mathbf{Rel}_\theta(x, y)$ the set of relevant elements. Note that, for all sufficiently large θ , the distance formula implies that $\mathbf{Rel}_\theta(x, y)$ is finite. In fact, using Lemma 2.5 of [BHS19], one can bound its cardinality in terms of θ, E , and $d_{\mathcal{X}}(x, y)$ without using the distance formula.

The notion of the rank of $(\mathcal{X}, \mathfrak{S})$ is easy to define, but it is of significant importance in the present paper:

Definition 1.10 (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 $U_1, \dots, U_n \in \mathfrak{S}$ for which $\pi_{U_i}(\mathcal{X})$ is unbounded for all i .

The rank is closely related to *standard product regions* in \mathcal{X} , which are a useful tool whose construction we now review; see also [BHS17b, Section 13] and [BHS19, Section 5]. 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 1.11 (Nested partial tuples). Recall that $\mathfrak{S}_U = \{V \in \mathfrak{S} \mid V \sqsubseteq U\}$. Fix $\kappa \geq E$ and let F_U be the set of κ -consistent tuples in $\prod_{V \in \mathfrak{S}_U} 2^{C^V}$.

Definition 1.12 (Orthogonal partial tuples). Let $\mathfrak{S}_U^\perp = \{V \in \mathfrak{S} \mid V \perp U\}$. Fix $\kappa \geq E$ and let E_U be the set of κ -consistent tuples in $\prod_{V \in \mathfrak{S}_U^\perp} 2^{C^V}$.

Definition 1.13 (Standard product regions in \mathcal{X}). Given \mathcal{X} and $U \in \mathfrak{S}$, there are coarsely well-defined maps $\phi^\sqsubseteq, \phi^\perp: F_U, E_U \rightarrow \mathcal{X}$ which extend to a coarsely well-defined map $\phi_U: F_U \times E_U \rightarrow \mathcal{X}$. 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 $(\phi_U(\vec{a}, \vec{b}))_V = \rho_V^U$. Finally, if $U \sqsubset V$, let $(\phi_U(\vec{a}, \vec{b}))_V = \rho_V^U$. We refer to $F_U \times E_U$ as a *standard product region*, whose image in \mathcal{X} we also call a *standard product region* and denote by P_U .

The image of F_U in \mathcal{X} is again a hierarchically hyperbolic space, with index set \mathfrak{S}_U and hyperbolic spaces and projections inherited from those in \mathfrak{S} . The same is true for E_U ,

although one must replace \mathfrak{S}_U with the set of $V \in \mathfrak{S}$ such that $V \perp U$, together with some element $A \in \mathfrak{S}$ into which each such V is nested (such an A is provided by the HHS axioms). We won't have much need for this here, and refer the interested reader to [BHS19, Section 5] for details.

1.1.2. *Hierarchically hyperbolic groups.* A finitely generated group G is a *hierarchically hyperbolic group (HHG)* if some (hence any) Cayley graph of G is an HHS, and the HHS structure is G -invariant. Specifically, an HHG is a finitely generated group G , equipped with a specific word metric, so that there is an HHS (G, \mathfrak{S}) where:

- G acts cofinitely on \mathfrak{S} , preserving each relation $\sqsubseteq, \perp, \triangleleft$;
- for each $U \in \mathfrak{S}$ and $g \in G$, there is an isometry $g: \mathcal{C}U \rightarrow \mathcal{C}gU$, and if $h \in G$, then the isometry $gh: \mathcal{C}U \rightarrow \mathcal{C}ghU$ is the same as the composition $\mathcal{C}U \xrightarrow{h} \mathcal{C}hU \xrightarrow{g} \mathcal{C}ghU$;
- for each $U \in \mathfrak{S}$ and $g, x \in G$, the points $g \circ \pi_U(x)$ and $\pi_{gU}(gx)$ uniformly coarsely coincide;
- for each $U, V \in \mathfrak{S}$ such that $U \triangleleft V$ or $U \sqsupseteq V$, and each $g \in G$, we have $\rho_{gU}^g = g(\rho_V^U)$.

Examples of hierarchically hyperbolic groups include mapping class groups of finite-type orientable surfaces and fundamental groups of compact special cube complexes, see [BHS17b, BHS19] for details and additional examples.

The only property of HHGs that we use in this paper is immediate from the definition, in particular from the property that G acts cofinitely on \mathfrak{S} : there exists $C \geq 0$ such that for all $U \in \mathfrak{S}$, either $\text{diam}(\mathcal{C}U) \leq C$, or $\mathcal{C}U$ has unbounded diameter.

1.1.3. *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 HHGs; it is also satisfied for all naturally occurring examples of HHSs. We impose it in order to rule out product regions with bounded but arbitrarily large factors. Our results likely have analogues that hold in the absence of this hypothesis, but would require custom-tailoring to the situation at hand.

Definition 1.14 (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 HHSs:

Theorem 1.15 (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 $f: (B_R(0) \subseteq \mathbb{R}^n) \rightarrow \mathcal{X}$ for all $R \geq 0$. 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 \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 coarsely decomposing as product whose factors are various subspaces F_U , so that any ultralimit of standard boxes in \mathcal{X} is homeomorphic to a subset of \mathbb{R}^ν because \mathcal{X} is asymphoric. Hence, $n \leq \nu$, as required. \square

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

Definition 1.16 (Hull of a set; Section 6 of [BHS19]). For each $A \subset \mathcal{X}$ and $\theta \geq 0$, 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 $\text{hull}_{\mathcal{C}W}(A)$, the coarse 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 $A \subset H_\theta(A)$.

Hulls are examples of *hierarchically quasiconvex* subspaces of \mathcal{X} . The other notable examples are standard product regions. The idea behind hierarchical quasiconvexity is to simultaneously capture (in our coarse setting) various notions of (coarse) convexity:

- First, hierarchical quasiconvexity directly generalizes the usual notion of quasiconvexity in a hyperbolic space: when \mathcal{X} is a hyperbolic HHS, the two notions coincide. More generally, hierarchical quasiconvexity of a subspace $\mathcal{Y} \subset \mathcal{X}$ requires that \mathcal{Y} has uniformly quasiconvex projections to all hyperbolic spaces $\mathcal{C}U$ for $U \in \mathfrak{S}$.
- Second, hierarchical quasiconvexity imitates, in the HHS setting, the notion of a convex subcomplex A of a CAT(0) cube complex M . That notion has many equivalent formulations; one of them says that M is convex provided that the median of x, y, z lies in A whenever at least two of the vertices x, y, z lie in A . This generalizes naturally to a notion of coarse median convexity in Bowditch’s *coarse median spaces* [Bow13], which are discussed in more detail below. It was verified in [BHS19, Section 7] that HHSs are coarse median spaces (we rely heavily on this fact in the rest of the paper) and that hierarchically quasiconvex subspaces are coarsely median convex. Recently, Russell-Spriano-Tran have proved the converse [RST18].
- From a point of view that emphasizes paths rather than coarse medians or projections, hierarchically quasiconvex subspaces are “quasiconvex with respect to hierarchy paths”: if \mathcal{Y} is hierarchically quasiconvex, then any hierarchy path with endpoints in \mathcal{Y} stays close to \mathcal{Y} .

A subset $\mathcal{Y} \subset \mathcal{X}$ is *hierarchically quasiconvex* if it has quasiconvex projections to the various hyperbolic spaces, and coarsely contains all realization points for tuples whose U -coordinate lies in $\pi_U(\mathcal{Y})$ for all $U \in \mathfrak{S}$. More precisely:

Definition 1.17 (Hierarchical quasiconvexity, Definition 5.1 of [BHS19]). Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then $\mathcal{Y} \subseteq \mathcal{X}$ is *k -hierarchically quasiconvex*, for some $k: [0, \infty) \rightarrow [0, \infty)$, 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 $\kappa \geq 0$ 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 $d_U(\pi_U(x), b_U) \leq \theta_e(\kappa)$ (where $\theta_e(\kappa)$ is as in Theorem 1.2) satisfies $d(x, \mathcal{Y}) \leq k(\kappa)$.

As one might expect, hulls of arbitrary sets are hierarchically quasiconvex, although in this paper we mainly consider hulls of finite sets:

Proposition 1.18. [BHS19, Lemma 6.2] *There exists θ_0 so that for each $\theta \geq \theta_0$ there exists $\kappa: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ so that for each $A \subset \mathcal{X}$ the set $H_\theta(A)$ is κ -hierarchically quasiconvex.*

Remark 1.19. Whenever we are working with a fixed HHS $(\mathcal{X}, \mathfrak{S})$, the notation θ_0 will refer to the constant from Proposition 1.18, and we fix once and for all a constant $\theta \geq \theta_0$.

1.2.1. *The gate map to a hierarchically quasiconvex subspace, and the bridge lemma.* We now recall a construction from Section 5 of [BHS19], namely the *gate map* to a hierarchically quasiconvex subspace, and prove some additional facts about it. (The terminology is inspired

by the similarity with the notion of a gate map to a convex subspace of a median space; see Section 1.5.)

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 $p_{U,A}: \mathcal{C}U \rightarrow \pi_U(A)$. Define a *gate map* $\mathfrak{g}_A: \mathcal{X} \rightarrow A$ as follows: given $x \in \mathcal{X}$, for each $U \in \mathfrak{S}$ let $b_U = p_{U,A}(x)$. In [BHS19, Section 5] we show that the tuple $(b_U)_{U \in \mathfrak{S}}$ is uniformly (depending on $\kappa(0)$) consistent, so Theorem 1.2 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 $U \in \mathfrak{S}$.

Intuitively, the gate map \mathfrak{g}_A takes x to some realization point for the tuple whose U -coordinate, for each U , is a closest point to $\pi_U(x)$ in the quasiconvex subspace $\pi_U(A)$.

The following lemma, Lemma 1.20 (“the bridge 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. Before we state it, we give some intuition coming from CAT(0) cube complexes:

Remark. The bridge lemma is well-illustrated by an analogy to CAT(0) cube complexes, where the notion was introduced by Behrstock–Charney [BC11]. In the analogy, let P, Q be convex subcomplexes of a CAT(0) cube complex \mathcal{M} , and let $\mathfrak{g}_P, \mathfrak{g}_Q: \mathcal{M} \rightarrow P, Q$ be cubical closest-point projection; on the 0-skeleton, these are the usual gate maps in the median space sense. (So, a hyperplane separates $\mathfrak{g}_P(x)$ from $\mathfrak{g}_P(y)$ if and only if it crosses P and separates x, y .) Then $\mathfrak{g}_P(Q)$ is a convex subcomplex of P crossed by exactly those hyperplanes that cross P and Q , and $\mathfrak{g}_Q(P)$ is a convex subcomplex of Q crossed by the same hyperplanes. The convex hull of $\mathfrak{g}_P(Q) \cup \mathfrak{g}_Q(P)$ is crossed by the above hyperplanes, together with the hyperplanes that separate P from Q . Hyperplanes of the latter type cross hyperplanes of the former type, and so the convex hull decomposes as a product, which one can view as a “bridge” between P and Q , in the sense that combinatorial geodesics from P to Q travel through the bridge.

Lemma 1.20 is analogous, except we have replaced the CAT(0) cube complex with an HHS, replaced convexity with hierarchical quasiconvexity, and replaced the cubical closest-point projection with the gate map.

Lemma 1.20 will be important later on in the paper. We use it in Section 3 to study boxes in asymptotic cones of an HHS; we use it in Section 4 to study coarse intersections between standard orthants, the key point being that if A, B are hierarchically quasiconvex, then the “coarse intersection” of A and B coarsely coincides with $\mathfrak{g}_A(B)$. Also, this lemma is useful for obtaining simplifications of the distance formula in various instances, see for instance Corollary 1.28 where we obtain a formula for the distance between a point and a product region. We note that another inspiration for this lemma is its analogue in the mapping class group, as developed in [BKMM12, Section 3].

Lemma 1.20 (Bridge lemma). *For every $\kappa: [0, \infty) \rightarrow [0, \infty)$ and all $K_0 \geq 10\kappa(0)E$, the following holds. There exists a function κ' and constants $K_1 = K_1(\kappa, E, K_0)$ and $K_2 = K_2(\kappa, E, K_0)$ and $K_3 = K_3(\kappa, E, K_1)$ such that for all κ -hierarchically quasiconvex sets A, B , we have:*

- (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 $\mathfrak{g}_A(B) \rightarrow \mathfrak{g}_A(B)$.

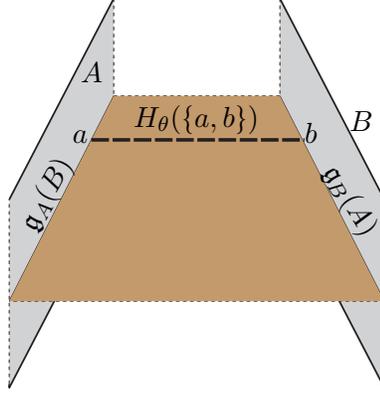


FIGURE 2. The bridge between quasi-convex sets A and B

- (3) For any $a \in \mathfrak{g}_A(B), b = \mathfrak{g}_B(a)$, we have a (K_1, K_1) -quasi-isometric embedding $f: \mathfrak{g}_A(B) \times H_\theta(\{a, b\}) \rightarrow \mathcal{X}$ with image $H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$, so that $f(\mathfrak{g}_A(B) \times \{b\})$ K_1 -coarsely coincides with $\mathfrak{g}_B(A)$.

Let $\mathcal{H} = \{U \in \mathfrak{S} : \text{diam}(\mathfrak{g}_A(B)) > K_0\}$. Let a, b, f be as above.

- (4) For each $p, q \in \mathfrak{g}_A(B)$ and $t \in H_\theta(\{a, b\})$, we have $\mathbf{Rel}_{K_2}(f(p, t), f(q, t)) \subseteq \mathcal{H}$.
 (5) For each $p \in \mathfrak{g}_A(B)$ and $t_1, t_2 \in H_\theta(\{a, b\})$, we have $\mathbf{Rel}_{K_2}(f(p, t_1), f(p, t_2)) \subseteq \mathcal{H}^\perp$.
 (6) For each $p \in A, q \in B$ we have

$$d(p, q) \asymp_{K_3, K_3} d(p, \mathfrak{g}_A(B)) + d(q, \mathfrak{g}_B(A)) + d(A, B) + d(\mathfrak{g}_{\mathfrak{g}_B(A)}(p), \mathfrak{g}_{\mathfrak{g}_B(A)}(q)).$$

The reader is referred to Figure 2 for a heuristic picture of the content of the lemma.

Proof of Lemma 1.20. 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 $d_V(A, B) \geq 100E\kappa(0)$. Fix $K_0 \geq 10E\kappa(0)$ and let \mathcal{H}_{K_0} be the set of $H \in \mathfrak{S}$ with $d_H(a, a') > K_0$ for some $a, a' \in \mathfrak{g}_A(B)$, say $a = \mathfrak{g}_A(b), a' = \mathfrak{g}_A(b')$ for some $b, b' \in B$.

The following claim can be proved using standard thin quadrilateral arguments in the hyperbolic space \mathcal{CV} for each $V \in \mathcal{V}$:

Claim 1.21. $\pi_V(\mathfrak{g}_A(B))$ and $\pi_V(\mathfrak{g}_B(A))$ have diameter $\leq 10E\kappa(0)$ for $V \in \mathcal{V}$.

For $U \in \mathfrak{S} - \mathcal{V}$ and $x \in \mathfrak{g}_A(B)$, $d_U(x, \mathfrak{g}_B(x)) \leq 10E\kappa(0)$.

The next auxiliary claim is a sufficient condition for orthogonality between $H \in \mathcal{H}_{K_0}$ and $V \in \mathcal{V}$:

Claim 1.22. Let $C \geq E$ and let $a, b, a', b' \in \mathcal{X}$ and suppose that $H, V \in \mathfrak{S}$ satisfy

- $d_V(a, a'), d_V(b, b') \leq C$;
- $d_V(a, b) > 10C$;
- $d_H(a, b), d_H(a', b') \leq C$;
- $d_H(a, a') > 10C$;

Then $H \perp V$.

Proof of Claim 1.22. To establish that $H \perp V$ we must show that H and V are not related by either the transversality or the nesting relation. Our proof is by contradiction.

Suppose $V \triangleleft H$. First, assume that we are in the case that $d_V(a, \rho_V^H) \leq E$. We then have that $d_V(\rho_V^H, b) > 8C$ and thus $d_V(\rho_V^H, b') > 6C$. Then, by consistency ρ_V^H lies E -close to both $\pi_H(b), \pi_H(b')$, which is impossible since $d_H(b, b') > 6C$. It remains to consider the case where $d_V(a, \rho_V^H) > E$. Here, by consistency, we have that $d_H(a, \rho_V^H) \leq E$. Hence $d_H(a', \rho_V^H) \geq 5E$,

and so, by consistency, we have $d_V(a', \rho_V^H) \leq E$. This case now reduces to the first case, with a' replacing a , again yielding a contradiction.

Suppose $V \not\subseteq H$. Since $d_H(a, a') > 10C$ and $d_H(b, b') > 6C$, at least one of the pairs a, b or a', b' has the property that geodesics in \mathcal{CH} connecting the corresponding projection points are E -far from ρ_H^V . By the bounded geodesic image axiom, we have, say, $d_V(a, b) \leq E$, a contradiction. The same argument rules out $H \not\subseteq V$.

Since we have ruled out nesting and transversality, we thus have $H \perp V$. \square

The preceding two claims imply that $V \perp H$ for all $V \in \mathcal{V}$ and $H \in \mathcal{H}_{K_0}$. We now proceed to the proofs of the enumerated assertions.

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 1.2 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 1.21. This proves Assertion (1).

For $b \in B$, Claim 1.21 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 $100E\kappa(0) \leq d_V(a, b) \leq d_V(A, B) + 20E\kappa(0)$ for $V \in \mathcal{V}$; the second inequality here follows from Claim 1.21. Since $a \in A$ and $b \in B$ we also have $d_V(A, B) \leq d_V(a, b)$.

Let $a' \in \mathfrak{g}_A(B)$. Assertion 2 implies that, up to uniformly bounded distance, $a' = \mathfrak{g}_A(b')$ for some $b' \in \mathfrak{g}_B(A)$. For each $U \in \mathfrak{S} - \mathcal{V}$, 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 1.23. *For each a', h as above, the associated tuple $(b_W)_{W \in \mathfrak{S}}$ defined above is $20K_0$ -consistent.*

Proof of Claim 1.23. 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 $\mathcal{CW}, \mathcal{CW}'$ of a common point in those cases.

If $W \in \mathfrak{S} - \mathcal{V}$ and $V \in \mathcal{V}$, then either

- $W \in \mathcal{H}_{K_0}$, or
- $\text{diam}_W(\pi_W(\mathfrak{g}_A(B))) \leq K_0$ and $d_W(a, b) \leq 100E\kappa(0)$.

In the first case, $V \perp W$ by Claim 1.22, 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 $W \not\perp V$, 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_V^W , 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. \square

Assertion (3): Given the consistent tuple provided by Claim 1.23, the realization theorem, Theorem 1.2, then provides a coarsely unique $x \in \mathcal{X}$ realizing (b_W) , and we let $f(a', h) = x$. This gives a map $f: \mathfrak{g}_A(B) \times H_\theta(a, b) \rightarrow \mathcal{X}$, and one can see using the distance formula that there exists $K_1 = K_1(\kappa, E)$ so that f is a (K_1, K_1) -quasi-isometric embedding. In the next claims, we check that f satisfies the remaining properties of Assertion (3).

Claim 1.24. $f(\mathfrak{g}_A(B) \times H_\theta(\{a, b\}))$ is coarsely contained in $H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$.

Proof of Claim 1.24. Let $h \in H_\theta(\{a, b\})$. Let $c \in B$ and let $x = f(\mathfrak{g}_A(c), h)$. 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 the geodesic γ_U in \mathcal{CU} from $\pi_U(a)$ to $\pi_U(b)$, by the definition of a θ -hull.

If $U \in \mathfrak{S} - \mathcal{V}$, then $\pi_U(x)$ lies uniformly close to $\pi_U(\mathfrak{g}_A(c))$. 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 $H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$. \square

Claim 1.25. $H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$ is coarsely contained in the image of f .

Proof of Claim 1.25. 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 1.21, 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. \square

Claim 1.26. $\mathfrak{g}_B(A)$ coarsely coincides with $f(\mathfrak{g}_A(B) \times \{b\})$.

Proof of Claim 1.26. By Claim 1.25, $\mathfrak{g}_B(A)$ 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 1.21. 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 $d_H(A, B) \leq 100E\kappa(0)$, 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 $\mathfrak{g}_B(A)$. \square

Assertions (4),(5): Let $p, q \in \mathfrak{g}_A(B)$ and $t_1, t_2 \in H_\theta(\{a, b\})$. Then there exists K_2 , depending on κ, K_1, E such that the following hold by the construction of f . First, if $H \in \mathbf{Rel}_{K_2}(f(p, t_1), f(q, t_1))$, then $H \in \mathcal{H}_{K_0}$.

Second, if $V \in \mathbf{Rel}_{K_2}(f(p, t_1), f(p, t_2))$, then $V \in \mathcal{V}$, so $V \in \mathcal{H}_{K_0}^\perp$ by Lemma 1.22, 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 1.27 provides $K_4 = K_4(\kappa, \mathcal{X})$ so that

$$d(\mathfrak{g}_F(p), \mathfrak{g}_F(q)) \asymp_{K_4, K_4} d(A, B) + d(\mathfrak{g}_{\mathfrak{g}_B(A)}(p), \mathfrak{g}_{\mathfrak{g}_B(A)}(q)),$$

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

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

for some uniform T . Given $L \geq 0$, let $\sigma_L(p, q) = \sum_{U \in \mathfrak{S}} \{\{d_U(p, q)\}\}_L$.

By the distance formula (Theorem 1.3), $d(p, q) \geq K_3^{-1} \sigma_{10T}(p, q) - K_3$ for some K_3 . Since, $10\sigma_{10T}(p, q) \geq \sigma_{100T}(p, \mathfrak{g}_F(p)) + \sigma_{100T}(\mathfrak{g}_F(p), \mathfrak{g}_F(q)) + \sigma_{100T}(\mathfrak{g}_F(p), q)$, the claim follows from another use of the distance formula (on the right, with threshold $100T$). \square

The next lemma is used in the proof of the final assertion of Lemma 1.20, but it is also interesting in its own right, since it says in particular that $\mathfrak{g}_A(a)$ is the ‘‘coarsely closest point’’ of the hierarchically quasiconvex set A to the (arbitrary) point $a \in \mathcal{X}$.

Lemma 1.27. *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 $d(a, B) \asymp_{K,K} d(a, \mathfrak{g}_B(a))$. Moreover, for any $a \in A$:*

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

Proof. First let $a \in \mathcal{X}$ and $b \in B$. Recall that for $U \in \mathfrak{S}$, the map $p_{U,B}: \mathcal{CU} \rightarrow \pi_U(B)$ is coarsely the closest-point projection. For any $U \in \mathfrak{S}$, we have $d_U(a, p_{U,B}(\pi_U(a))) \leq d_U(a, b) + 1$. By the definition of the gate, and the distance formula, we thus have K' , depending on κ , so that $d(a, \mathfrak{g}_B(a)) \leq K'd(a, b) + K'$. 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{CU} -geodesic from $\pi_U(a)$ to $p_{U,B}(\pi_U(a))$, so by the distance formula and the definition of the gate, $d(a, \mathfrak{g}_B(a)) \geq d(\mathfrak{g}_B(a), \mathfrak{g}_A(\mathfrak{g}_B(a)))/K' - K'$ for K' depending only on $\mathcal{X}, \mathfrak{S}$, and κ .

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

Although we will not use it in the rest of the paper, we note the following interesting corollary, which is useful elsewhere:

Corollary 1.28. *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. Then for all sufficiently large s , there exists K such that the following holds. Let $U \in \mathfrak{S}$. Let P_U be a corresponding standard product region and let $x \in \mathcal{X}$. Let \mathcal{R} be the set of $V \in \mathfrak{S}$ such that $U \sqsubset V$ or $U \triangleleft V$ and $d_V(x, \rho_V^U) > s$. Then*

$$d(x, P_U) \asymp_{K,K} \sum_{V \in \mathcal{R}} d_V(x, \rho_V^U).$$

Proof. By construction, P_U is κ -hierarchically quasiconvex, where κ depends only on E . Lemma 1.27 provides K_0 such that $d(x, P_U) \asymp_{K_0} d(x, \mathfrak{g}_{P_U}(x))$. Now, by the definition of \mathfrak{g}_{P_U} , the projections $\pi_V(x)$ and $\pi_V(\mathfrak{g}_{P_U}(x))$ uniformly coarsely coincide unless $U \sqsubset V$ or $U \triangleleft V$. In the latter case, $\mathfrak{g}_{P_U}(x)$ projects uniformly close to ρ_V^U , by Definition 1.11 and Definition 1.12. The claim now follows from the distance formula, Theorem 1.3. \square

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

Definition 1.29 (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 will always be the case in this paper, any orientation is canonical.

Definition 1.30 (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, as was proven independently by Chatterji-Niblo [CN05] and Nica [Nic04], building on work of Sageev [Sag95]. Thus this graph is the 1-skeleton of a uniquely determined CAT(0) cube complex, by a theorem of Chepoi [Che00]; we call this complex 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 CAT(0) cube complex dual to the resulting wallspace is exactly C .

Definition 1.31 (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 .

Finally, we need the notion of a *cubical orthant*. Let C be a CAT(0) cube complex. Let $n \geq 1$ and let \mathbf{R} be the standard tiling of $[0, \infty)$ by 1-cubes. A *cubical n -orthant* is a copy of the CAT(0) cube complex \mathbf{R}^n with the obvious product cubical structure. A cubical n -orthant *in* C is a subcomplex O of C that is isomorphic to \mathbf{R}^n and has the property that $O \hookrightarrow C$ is an isometric embedding (and in particular a median homomorphism) on the 0-skeleton.

1.4. Ultralimits and asymptotic cones. We now recall the definitions of ultralimits and asymptotic cones of metric spaces. A more detailed discussion can be found, for example, in the book [DK18] or in [Dru02]; we recall just the notions we need.

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.,

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

Then \mathbf{M} is the associated quotient metric space, obtained from \widehat{M} by identifying points y and z for which $\mathbf{d}(y, z) = 0$.

More generally, given a sequence (M_n, d_n) of metric spaces, with a basepoint $m_n \in M_n$ for each n , we define the ultralimit as follows. Given $x = (x_n), y = (y_n) \in \prod_n M_n$, let $\mathbf{d}(x, y) = \lim_{\omega} d_n(x_n, y_n)$. We identify $(x_n), (y_n)$ when $\mathbf{d}(x, y) = 0$, and restrict ourselves to points (x_n)

for which $\lim_{\omega} d(x_n, m_n) < \infty$. The resulting based space is the ultralimit $\lim_{\omega} (M_n, d_n, m_n)$. Note that the asymptotic cone \mathbf{M} defined above is just $\lim_{\omega} (M_n, d_n/s_n, m_n)$. When taking ultralimits of a sequence of spaces without rescaling, we will emphasize this by saying “non-rescaled”.

We will adopt the following notational conventions for asymptotic cones. We let ω denote a non-principal ultrafilter on \mathbb{N} , fixed once and for all. 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, quasimedial. We recall some background on median spaces and coarse median spaces. The latter were introduced by Bowditch [Bow13] and we refer the reader to [Bow13, Bow18b] for a more detailed discussion of both concepts.

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 $\mu: (\mathcal{Y}^{(0)})^3 \rightarrow \mathcal{Y}^{(0)}$ 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 1.32 (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 1.33 (Coarse median space; [Bow13]). Let (\mathcal{L}, d) be a metric space and let $\mu: \mathcal{L}^3 \rightarrow \mathcal{L}$ 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 $h: \mathbb{N} \rightarrow [0, \infty)$ so that the following hold:

- For all $x, y, z, x', y', z' \in \mathcal{L}$,

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

It was shown in [BHS19] that every hierarchically hyperbolic space is a coarse median space; we refer the reader there for details of the construction. The property of coarse medians we need in this paper is that, given an HHS $(\mathcal{X}, \mathfrak{S})$, there exists a constant ℓ , depending only on the HHS constant E , so that the following holds. Given $x, y, z \in \mathcal{X}$ and letting $m \in \mathcal{X}$ be their coarse median, then for all $U \in \mathfrak{S}$, the point $\pi_U(m)$ lies ℓ –close to any geodesic in \mathcal{CU} joining a, b , where $a, b \in \{\pi_U(x), \pi_U(y), \pi_U(z)\}$ are distinct.

Definition 1.34 (Quasimedial map). Let \mathcal{Y} be a CAT(0) cube complex with median map $\mu_{\mathcal{Y}}$ on its 0-skeleton. Let (\mathcal{L}, μ, d) be a coarse median space. Let $h \geq 0$. An h -quasimedial map is a map $q: \mathcal{Y} \rightarrow \mathcal{L}$ for which

$$d(\mu(q(x), q(y), q(z)), q(\mu_{\mathcal{Y}}(x, y, z))) \leq h$$

for all $x, y, z \in \mathcal{Y}$.

Note that quasimedial maps are referred to by [Bow13] as “quasimorphisms,” but we use a different terminology to avoid any confusion with other uses of that word.

When studying asymptotic cones of HHSs, it isn’t sufficient to restrict oneself to finite median algebras/CAT(0) cube complexes. So, we need a few more standard notions about general median algebras and median metric spaces.

Recall that a set \mathcal{M} equipped with a ternary operation $\mu: \mathcal{M}^3 \rightarrow \mathcal{M}$ 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 1.33 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}, μ) 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 $d(a, b) = d(a, c) + d(c, b)$. 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.

Bowditch showed, in [Bow13, Theorem 2.3], that any asymptotic cone of a coarse median space of rank ν is a topological median algebra 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 . Moreover, Bowditch showed in [Bow18b, Proposition 2.4] (see also Theorem 6.9 of the same paper) that any asymptotic cone of a coarse median space is bilipschitz homeomorphic to a metric median space, where the median is as just described. When we work with asymptotic cones of HHSs (recall that each HHS is coarse median of finite rank), we will only be interested in the bilipschitz homeomorphism class, and will therefore assume that the asymptotic cone, with the given median, is a median metric space.

We collect the above in the following proposition, which plays an important role throughout Section 3:

Proposition 1.35 (Asymptotic cones of HHS are median metric spaces). *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Let \mathcal{X} be an asymptotic cone of \mathcal{X} . Let $\mu: \mathcal{X}^3 \rightarrow \mathcal{X}$ be the coarse median map, and let $\boldsymbol{\mu}: \mathcal{X}^3 \rightarrow \mathcal{X}$ be the map sending $\mathbf{x} = (x_n), \mathbf{y} = (y_n), \mathbf{z} = (z_n)$ to the point represented by $(\mu(x_n, y_n, z_n))$. Then:*

- $\boldsymbol{\mu}$ makes \mathcal{X} into a topological median space of finite rank. If $(\mathcal{X}, \mathfrak{S})$ is asymptotic and has rank ν , then the median space \mathcal{X} has rank at most ν .
- \mathcal{X} , equipped with the median $\boldsymbol{\mu}$, is bilipschitz equivalent to a median metric space.

Proof. It is shown in [BHS19, Section 7] that \mathcal{X} is a coarse median space. From Theorem F it follows that the rank of \mathcal{X} as a coarse median space is bounded above by the maximal cardinality m of collections $\{U_i\} \subset \mathfrak{S}$ of pairwise orthogonal elements. (So, in general, the coarse median rank of \mathcal{X} is bounded only by the complexity of \mathfrak{S} .)

The bound on the coarse median rank in the asymptotic case is Corollary 2.15 below.

The first assertion now follows from [Bow13, Theorem 2.3] which provides the median structure and implies that ν is an upper bound on the rank of the median space (\mathcal{X}, μ) .

The second assertion follows from [Bow18b, Theorem 6.9]. More specifically, being an asymptotic cone of a quasigeodesic space, \mathcal{X} is a complete geodesic space (see e.g., [DK18, Proposition 10.70]). The proof of [Bow13, Proposition 9.1] shows that \mathcal{X} , with the given median, satisfies the hypotheses of [Bow18b, Proposition 2.4], which then yields the second assertion. \square

Finally, we conclude with some background about the notion of gate maps in a median space, and the notion of a block; these again are vital in Section 3.

Definition 1.36 (Block, median gate). Let (M, μ, 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 in \mathbb{R} , endowed with the ℓ_1 metric.

Recall that the (median) interval in M between points m and n is the set $[m, n]$ of all m' such that $\mu(m, m', n) = m'$.

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

Closed convex subsets of a complete median space always admit a unique gate map (see e.g. [DK18, Lemma 6.26]). If N, N' are median convex, then $\mathfrak{g}_N(N')$ is again median convex; see [Bow18b].

1.6. Identifying hierarchy paths. We now prove a sufficient condition for a path in the HHS $(\mathcal{X}, \mathfrak{S})$ to be a hierarchy ray. It is straightforward, but it will play a role in Section 4.

In the lemma, “quasimedial” will mean with respect to the coarse median on \mathcal{X} and the usual median on \mathbb{R} , i.e., $\gamma: \mathbb{R} \rightarrow \mathcal{X}$ is quasimedial if whenever $r, s, t \in \mathbb{R}$ satisfy $r < s < t$, then the coarse median of $\gamma(r), \gamma(s), \gamma(t)$ is λ -close to $\gamma(s)$.

Lemma 1.37. *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. Then for all $\lambda \geq 1$, there exists $D = D(\lambda)$ such that the following holds. Let $I \subset \mathbb{R}$ be a (possibly unbounded) subinterval and let $\gamma: I \rightarrow \mathcal{X}$ be a λ -quasimedial (λ, λ) -quasi-isometric embedding. Then γ is a (D, D) -hierarchy path.*

Proof. The path γ is a (λ, λ) -quasigeodesic by hypothesis, so to show that it is a hierarchy path we only need to prove that there exists a constant D so that, for each $U \in \mathfrak{S}$, the composition of γ with π_U is an unparameterized (D, D) -quasigeodesic in $\mathcal{C}U$. In order to do so, it suffices to show that there exists a constant D' so that for each $r, s, t \in I$ with $r < s < t$, we have that $\pi_U(\gamma(s))$ lies D' -close to a geodesic from $\pi_U(\gamma(r))$ to $\pi_U(\gamma(t))$.

Let $r, s, t \in I$ satisfy $r < s < t$. Let m be the coarse median of $\gamma(r), \gamma(s), \gamma(t)$. Since π_U is E -coarsely Lipschitz and γ is λ -quasimedial, we have $d_U(\pi_U(\gamma(s)), \pi_U(m)) \leq E\lambda + E$. By the definition of the coarse median, there exists $\lambda' = \lambda'(E, \lambda)$ such that $d_U(\gamma(s), m_U) \leq \lambda'$, where m_U is the coarse median in the hyperbolic space $\mathcal{C}U$ of the three points $\pi_U(\gamma(r)), \pi_U(\gamma(s)), \pi_U(\gamma(t))$. The distance from m_U to any geodesic $[\pi_U(\gamma(r)), \pi_U(\gamma(t))]$ is bounded in terms of the hyperbolicity constant of $\mathcal{C}U$, so we are done. \square

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, i.e., approximated by a finite CAT(0) cube complex in such a way that both distances and (coarse) medians are coarsely preserved.

We achieve the cubulation of $H_\theta(A)$ by constructing finitely many walls in $H_\theta(A)$ and then passing to the dual cube complex, using the work of Chatterji–Niblo, Nica, and Sageev mentioned in Definition 1.30.

In the case where $(\mathcal{X}, \mathfrak{S})$ is a rank-one HHS — which, as we will see below, is equivalent to being hyperbolic — the cubulation of the hull of A reduces to a classical fact about hyperbolic

spaces: the coarse convex hull of any finite collection of points can be approximated by a quasi-isometrically embedded 1–dimensional $\text{CAT}(0)$ cube complex, i.e., a tree [Gro87b].

We exploit this special case to build walls in $H_\theta(A)$, roughly as follows. We consider $U \in \mathfrak{S}$ and consider a tree which approximates the coarse convex hull of $\pi_U(A)$ in the hyperbolic space \mathcal{CU} . 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.

The fact that the quality of the tree approximation in \mathcal{CU} depends on $|A|$ is the most obvious way in which the dependence of the quality of our cubical approximation on $|A|$ makes itself felt. However, there are also several other essentially different ways in which $|A|$ influences the quality of the approximation. First, it does so in our choice of separated nets (roughly, the larger the total branching of a tree is, the harder it is to approximate the tree with a separated net), and the other two are in Lemma 2.6 and Lemma 2.10.

We now turn to the formal statement of the cubulation of hulls theorem (which is Theorem F of the introduction):

Theorem 2.1 (Cubulation of hulls). *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 $M \geq M_0$ there is a constant C_1 so that for any $A \subset \mathcal{X}$ of cardinality $\leq k$, there is a C_1 –quasimedial (C_1, C_1) –quasi-isometry $p_A: \mathcal{Y} \rightarrow H_\theta(A)$, where \mathcal{Y} is a $\text{CAT}(0)$ cube complex.*

Moreover, let \mathcal{U} be the set of $U \in \mathfrak{S}$ so that $d_U(x, y) \geq M$ 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 $y_1, \dots, y_{k'} \in \mathcal{Y}$ so that $k' \leq k$ and \mathcal{Y} is equal to the convex hull in \mathcal{Y} of $\{y_1, \dots, y_{k'}\}$.

Remark. Since we posted an earlier version of this paper, Bowditch has given a new proof of this theorem under somewhat more general hypotheses (very similar to, but strictly weaker than, the HHS axioms); see Theorem 1.3 in [Bow18a].

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 $A = \{x_1, \dots, x_k\} \subseteq \mathcal{X}$.

2.1. The candidate finite $\text{CAT}(0)$ 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{CU} of diameter at most E ; for each j , choose $\ell_j^U \in \pi_U(x_j)$, to obtain k points $\ell_1^U, \dots, \ell_k^U \in \mathcal{CU}$. As shown by Gromov, there exists $C = C(k, \delta)$ so that there is a finite tree T_U and an embedding $T_U \hookrightarrow \mathcal{CU}$, sending edges to geodesics of \mathcal{CU} , such that:

- $d_U(p, q) \leq d_{T_U}(p, q) \leq d_U(p, q) + C$ for all $p, q \in T_U$;
- $\ell_j^U \in T_U$ for $1 \leq j \leq k$;
- each leaf of T_U lies in $\{\ell_1^U, \dots, \ell_k^U\}$.

This is the usual spanning tree of a finite subset of a hyperbolic space; see [Gro87b]. The given properties of T_U ensure that, up to increasing C uniformly, $d_{\text{Haus}}(T_U, \text{hull}_{\mathcal{CU}}(\pi_U(A))) \leq C$.

Our choice of T_U ensures that, for each $x_j \in A$, every leaf of T_U is contained in $\pi_U(x_j)$ for some $x_j \in A$ and each $\pi_U(x_j)$ contains a point of T_U .

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 set of all $U \in \mathfrak{S}$ with $\text{diam}(\pi_U(A)) \geq 100Mk$.

Let $\mathcal{U}_1 \subseteq \mathcal{U}$ be the set of \sqsupseteq –minimal elements of \mathcal{U} . Given \mathcal{U}_{n-1} , let $\mathcal{U}_n \subseteq \mathcal{U}$ be the set of all \sqsupseteq –minimal elements of $\mathcal{U} - \mathcal{U}_{n-1}$. Finite complexity ensures that there is some s so that $\bigcup_{n=1}^s \mathcal{U}_n = \mathcal{U}$. For each

$U \in \mathcal{U}$, let $\mathcal{U}^{\sqsupseteq, U} = \{V \in \mathcal{U} : V \sqsupseteq U\}$. For each $V \in \mathcal{U}^{\sqsupseteq, U}$, 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).

Remark (Finiteness of \mathcal{U}). For sufficiently large M (in terms of the threshold constant in the distance formula), Theorem 1.3 implies that \mathcal{U} is finite. One could also deduce this from the large link axiom (Definition 1.1.(7) in [BHS19]), avoiding use of the distance formula. Finiteness of \mathcal{U} is used below; it will ensure that the wallspace we construct has finitely many walls, so that the 0–cubes of \mathcal{Y} correspond bijectively to coherent orientations of the walls (recall Definition 1.30); in other words, we don’t have to worry about the “canonical orientation” condition from Definition 1.30, because we will be dealing with a finite wallspace.

We now proceed to the construction of the walls.

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 implies that the p_i^U together with the r_V^V provide a $10Mk$ –net which is M –separated):

- (1) $d_U(p_i^U, x_j) \geq M$,
- (2) $d_U(p_i^U, p_j^U) \geq M$,
- (3) $d_U(p_i^U, r_V^V) \geq M$ for each $V \in \mathcal{U}^{\square, U}$ (when $U \in \mathcal{U}_1$, there are no such V), and
- (4) each component of $T_U - \left(\{p_i^U\} \cup \{r_V^V\}_{V \in \mathcal{U}^{\square, U}} \right)$ has diameter at most $10Mk$ (when $U \in \mathcal{U}_1$, there are no such V , so the criterion is only about complements of the $\{p_i^U\}$).

The existence of such a net is justified as follows. Fix $U \in \mathcal{U}$. For each $x_j \in A$, choose $y_j \in T_U$ lying in $\pi_U(x_j)$.

For each $j \leq k$, let T_U^j be the subtree of T_U spanned by y_1, \dots, y_j . Consider the geodesic T_U^2 . Let a_1, \dots, a_ℓ be the points a_s on T_U^2 such that there is a (possibly trivial) geodesic in T_U that intersects T_U^2 at a_s and joins a_s to a point in $\{y_3, \dots, y_k\}$. Note that $\ell \leq k - 2$.

Note that T_U is the union of T_U^2 along with ℓ subtrees C_s , each of which intersects T_U^2 at a point a_s , $s \leq \ell$. Choose a (possibly empty) M –separated set of points p_i^U in T_U^2 so that each p_i^U is M –far from each a_s , and M –far from y_1, y_2 , and M –far from each $\rho_V^V, V \in \mathcal{U}^{\square, U}$ belonging to T_U^2 . Any collection that is maximal with these properties has the property that $\{p_i^U\} \cup \{\rho_V^V \in T_U^2\}$ is an $M(\ell + 2)$ –net in T_U^2 . If $k = 2$, then $\{a_s\} = \emptyset$ and we are done.

Otherwise, each tree C_s contains at most $k - 2$ of the points y_j , and exactly one of the points a_1, \dots, a_ℓ , namely a_s . So, by induction, C_s contains an M –separated collection of points $\{p_i^U(s)\}_i$ that are M far from any $\rho_V^V \in C_s$, and M –far from any $y_j \in C_s$, and M –far from a_s , such that $\{p_i^U(s)\}_i \cup \{\rho_V^V \in C_s\}$ is an $M(k - 1)$ –net in C_s . Observe that the set of p_i^U , together with the union over s of the $\{p_i^U(s)\}_i$, has the properties listed above.

(Since the points p_i^U are M –far from each y_j , they are $(M - E)$ –far from $\pi_U(x_j)$, and so we rename $M - E$ to M to see that the first property on the list holds for $\pi_U(x_j)$, not just y_j . With the renamed constant, we now have an $(M + E)k$ –net, and in particular a $10Mk$ –net, provided $M \geq E/9$. We assume this just to simplify computations later.)

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 and the number of points x_j). 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).

The cubes in \mathcal{Y} are closely related to the standard product regions in \mathcal{X} . Specifically, each cube corresponds to a collection of pairwise-crossing walls. Each wall was determined by some $U \in \mathcal{U}$, namely the U for which the halfspaces in the wall are preimages of complementary components of some p_i^U . We think of U as labeling the wall. Now, each edge of \mathcal{Y} is labeled by the same U that labels its dual wall. Below, we will map \mathcal{Y} to \mathcal{X} in such a way that an edge e labeled U has the property that its endpoints are sent to points in \mathcal{X} that project uniformly close on \mathcal{CV} to ρ_V^U whenever $U \sqsubset V$ or $U \pitchfork V$. In this sense, we will be mapping cubes to standard product regions.

2.2. Lemmas supporting consistency of certain tuples. The map $\mathcal{Y} \rightarrow H_\theta(A)$ will be constructed roughly as follows. For each 0-cube $p \in \mathcal{Y}$, we will construct a tuple $(b_V) \in \prod_V \mathcal{CV}$. In Lemma 2.7, we will verify that this tuple is consistent, and this will require the following technical lemmas, which are essentially just applications of consistency (Axiom 1.1) and bounded geodesic image (Lemma 1.5).

The content of the lemmas is the following. Given distinct, non-orthogonal $U \in \mathcal{U}, V \in \mathfrak{S}$, there are three possibilities: we can have $U \pitchfork V, U \sqsubset V$, or $V \sqsubset U$. In the first two cases, the coarse point ρ_V^U lies close to T_V in \mathcal{CV} . In the second case, for any $x \in T_U$ far from ρ_U^V , the coarse point $\rho_V^U(x)$ lies close to T_V .

Lemma 2.3 (ρ_V^U close to T_V , transverse case). *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 $\text{diam}_{\mathcal{CU}}(\pi_U(A)) \geq 100Mk \geq 100M > 10^3E$. Hence we can choose $x_i \in A$ so that $d_U(x_i, \rho_U^V) > E$. Consistency yields $d_V(x_i, \rho_V^U) \leq E$. Since $\pi_V(x_i)$ has diameter $\leq E$ and contains a point of T_V , we have $d_V(T_V, \rho_V^U) \leq 2E$. \square

Lemma 2.4 (ρ_V^U close to T_V , nested case). *For any $M > 10E$, the following holds. Let $U \in \mathcal{U}, V \in \mathfrak{S}$, with $U \sqsubset V$. Then $d_V(\rho_V^U, T_V) \leq 100EC$.*

Proof. Suppose that $d_V(\rho_V^U, T_V) > 100EC$. Then, since T_V C -coarsely coincides with $\text{hull}_{\mathcal{CV}}(A)$, and the latter is $5E$ -quasiconvex, we have that ρ_V^U 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{CU} of diameter at most E , i.e., $\pi_U(A)$ has diameter bounded by $10E$. This contradicts $U \in \mathcal{U}$, provided $M > 10E$. \square

Lemma 2.5 ($\rho_V^U(x)$ close to T_V). *For any $M > 10EC$ the following holds. Consider $U \in \mathcal{U}$ and any $V \in \mathfrak{S}$ with $V \sqsubset U$. Then for each $x \in T_U - \mathcal{N}_M(\rho_U^V)$ there exists $x_j \in A$ with $d_V(\rho_V^U(x), x_j) \leq 2E$ (in particular, $\rho_V^U(x)$ 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 $T_U - \mathcal{N}_{M/2}(\rho_U^V)$ as x . Geodesics from x to $\pi_U(x_j)$ thus stay E -far from ρ_U^V , so that the desired conclusion follows from bounded geodesic image (and consistency, which says $\text{diam}_V(\pi_V(x_j) \cup \rho_V^U(\pi_U(x_j))) \leq E$). \square

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 $\mathfrak{p}_A: \mathcal{Y} \rightarrow \mathcal{X}$, 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}_i^U provided by Definition 2.2.

For $U \in \mathcal{U}, V \in \mathfrak{S}$ and each p_i^U (which we recall gives a pair $\{\overleftarrow{W}_i^U, \overrightarrow{W}_i^U\}$), we can consider $\overline{W}_i(U) \in \{\overleftarrow{W}_i^U, \overrightarrow{W}_i^U\}$ 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 $\beta_V(\overleftarrow{W}_i(U))$, where, as above, β_V is the composition of projection to \mathcal{CV} and the closest point projection to T_V .

By the definition of a coherent orientation, for any U, i, U', i' , we have $\beta_V(\overleftarrow{W}_i(U)) \cap \beta_V(\overleftarrow{W}_{i'}(U')) \neq \emptyset$, whence $S_{U,i,V}(p) \cap S_{U',i',V}(p) \neq \emptyset$. The Helly property for trees thus ensures that $\bigcap_{U,i} S_{U,i,V}(p) \neq \emptyset$ for each $V \in \mathfrak{S}$, and we let $b_V = b_V(p) = \bigcap_{U,i} S_{U,i,V}(p)$. Lemma 2.6, below, proves that $\text{diam}(b_V)$ are uniformly bounded. Lemma 2.7, 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 1.2. Specifically, there exists $\xi = \xi(\eta, E)$ so that for all $U \in \mathfrak{S}$, we have $\mathfrak{d}_U(\pi_U(\mathfrak{p}_A(p)), b_U) \leq \xi$.

The image of \mathfrak{p}_A coarsely coincides with $H_\theta(A)$: Let $x \in H_\theta(A)$.

For each wall in $H_\theta(A)$ (the walls are those from Definition 2.2), choose the halfspace containing x ; there is exactly one such halfspace since a wall is, by construction, a partition of $H_\theta(A)$ into two halfspaces. Now, any two of the chosen halfspaces contain x , so by Definition 1.29, this orientation is coherent, and it is a canonical orientation simply because there are only finitely many walls.

So, this orientation of all walls determines a 0-cube $p \in \mathcal{Y}$, by Definition 1.30. Now, by construction, the tuple $(b_V(p))$ has the property that, for all $U \in \mathcal{U}$, we have $\beta_U(x) \in b_U(p)$. Since $\mathfrak{d}_U(\pi_U(x), \beta_U(x)) \leq \theta$, because $x \in H_\theta(A)$, we see that $\mathfrak{d}_U(x, b_U(p)) \leq \theta$. Now, $\mathfrak{d}_U(b_U(p), \mathfrak{p}_A(p)) \leq \xi$, so $\mathfrak{d}_U(x, \mathfrak{p}_A(p)) \leq \xi + \theta$ for all $U \in \mathcal{U}$. If $U \notin \mathcal{U}$, then $\mathfrak{d}_U(x, \mathfrak{p}_A(p)) \leq \theta + 100Mk$. So, by the uniqueness axiom for HHS (Definition 1.1.(9) in [BHS19]), or simply by Theorem 1.3, we have $\mathfrak{d}_{\mathcal{X}}(x, \mathfrak{p}_A(p)) \leq C'_1$, where $C'_1 = C'_1(M, k, \mathcal{X}, \theta)$. Hence $H_\theta(A)$ lies in a uniform neighborhood of $\text{im } \mathfrak{p}_A$.

On the other hand, if $p \in \mathcal{Y}$, then $\pi_U(\mathfrak{p}_A(p))$ lies uniformly close (in terms of ξ) to $\text{hull}(\pi_U(A))$ for all $U \in \mathfrak{S}$. The definition of hierarchical quasiconvexity, together with the fact that $H_\theta(A)$ is hierarchically quasiconvex, ensures that $\mathfrak{p}_A(p)$ lies uniformly close to $H_\theta(A)$, i.e., $\text{im } \mathfrak{p}_A$ lies in a uniform neighborhood of $H_\theta(A)$.

After enlarging C'_1 if necessary, we thus see that there exists $C'_1 = C'_1(M, k, \mathcal{X}, \theta)$ such that $H_\theta(A)$ and $\text{im } \mathfrak{p}_A$ lie at Hausdorff distance at most C'_1 .

Distance estimates: For $p \in \mathcal{Y}$, we say p_i^U is a *separator* for p if p_i^U separates $\beta_U(\mathfrak{p}_A(p))$ from $b_U(p)$. We call U the *support* of the separator. In Lemma 2.10 we produce a constant $T = T(M, k, \eta, \xi, \mathfrak{S})$, so that for each $p \in \mathcal{Y}$ there are at most T separators for p .

We first 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 .

Specifically, let $p, q \in \mathcal{Y}$. By the definition of distance in a CAT(0) cube complex, $\mathfrak{d}_{\mathcal{Y}}(p, q)$ is the number of walls separating p and q . Let \mathcal{L}_i^V be a wall separating p from q . Then, by the construction of the tuples $b_V(p)$ and $b_V(q)$, the subtrees $b_V(p)$ and $b_V(q)$ of T_V lie on opposite sides of the wall in T_V determined by p_i^V . Conversely, if $b_V(p)$ and $b_V(q)$ are separated by the partition of T_V determined by some p_i^V , then \mathcal{L}_i^V corresponds to a wall in \mathcal{Y} separating p from q .

Hence $\mathfrak{d}_{\mathcal{Y}}(p, q)$ is the sum of the numbers of p_i^V separating $b_V(p)$ from $b_V(q)$, as V varies. Now, \mathcal{L}_i^V separates $b_V(p)$ from $b_V(q)$ but fails to separate $\beta_V(\mathfrak{p}_A(p))$ from $\beta_V(\mathfrak{p}_A(q))$ only if \mathcal{L}_i^V is a separator for p or for q . Similarly, \mathcal{L}_i^V separates $\beta_V(\mathfrak{p}_A(p))$ from $\beta_V(\mathfrak{p}_A(q))$ but fails to separate $b_V(p)$, $b_V(q)$ only if \mathcal{L}_i^V is a separator for p or for q .

Lemma 2.10 shows that p has at most T separators and q has at most T separators. Let $Q(p, q)$ be the sum over all V of the number of p_i^V separating $\beta_V(\mathfrak{p}_A(p))$ from $\beta_V(\mathfrak{p}_A(q))$. The preceding discussion shows that $|\mathfrak{d}_{\mathcal{Y}}(p, q) - Q(p, q)| \leq 2T$.

Observe that: if, for some V , there exist distinct p_i^V, p_j^V separating $\beta_V(\mathfrak{p}_A(p))$ from $\beta_V(\mathfrak{p}_A(q))$, then V contributes to the distance formula sum between $\mathfrak{p}_A(q)$ and $\mathfrak{p}_A(p)$, at some fixed threshold L chosen in terms of E and M . Moreover, V also contributes to the

distance formula sum in the case where $\beta_V(\mathfrak{p}_A(p))$ and $\beta_V(\mathfrak{p}_A(q))$ are both $M/2$ -close to $\pi_V(A)$ and there exists at least one p_i^V separating $\beta_V(\mathfrak{p}_A(p))$ from $\beta_V(\mathfrak{p}_A(q))$.

Applying Lemma 2.8 and Lemma 2.9, we have

$$d_{\mathcal{X}}(\mathfrak{p}_A(p), \mathfrak{p}_A(q)) \asymp \sum_{U \in \mathfrak{S}} \{\{d_U(\mathfrak{p}_A(p), \mathfrak{p}_A(q))\}\}_L \geq Q(p, q) - 100EC\theta N,$$

where N is the constant from Lemma 2.9. Hence there exists $C_1'' = C_1''(M, \mathcal{X}, \mathfrak{S}, k)$ so that $d_{\mathcal{X}}(\mathfrak{p}_A(p), \mathfrak{p}_A(q)) \geq d_{\mathcal{Y}}(p, q)/C_1'' - C_1''$ 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 2.12 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 it 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 2.13 provides C_1'''' so that \mathfrak{p}_A is C_1'''' -quasimedial, so the proof is complete once we take $C_1 = \max\{C_1', C_1'', C_1''', C_1''''\}$. \square

2.3.1. *Lemmas supporting realization.* The two lemmas below are used to construct a point in \mathcal{X} via realization, given the tuple $(b_V(p)) = (b_V)$ associated to a 0-cube $p \in \mathcal{Y}$ (which we fix for the purposes of the next two lemmas). The first lemma shows that b_V is a uniformly bounded set in each \mathcal{CV} , and the second verifies that the tuple (b_V) is η -consistent (and bounds η).

The realization theorem (Theorem 1.2) then provides a point $\mathfrak{p}_A(p) \in \mathcal{X}$ that projects ξ -close to b_V in each \mathcal{CV} , where ξ just depends on E and η . This is how we defined the map $\mathfrak{p}_A : \mathcal{Y} \rightarrow \mathcal{X}$ in the proof of Theorem 2.1.

Lemma 2.6. *There exists $\tau = \tau(M, k) > 0$ (independent of V) so that $\text{diam}(b_V(p)) \leq \tau$ for all $p \in \mathcal{Y}$.*

Proof. Fix $p \in \mathcal{Y}$ and write $b_V = b_V(p)$.

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

There exists $\tau = \tau(M, k) \geq 50Mk(k-2)$ such that the following holds. Suppose that $x, y \in \mathcal{X}$ satisfy $d_V(x, y) > \tau$. Then there exists $\alpha \in \{p_i^V\}_i \cup \{r_V^W\}_{W \in \mathcal{U}_1 \cap V \in \mathcal{V}}$ so that α is $10M$ -far from $\beta_V(x), \beta_V(y)$ and from all points of T_V of valence larger than 2, and separates $\beta_V(x)$ from $\beta_V(y)$. Indeed, there are at most $k-2$ points of valence larger than 2, since each leaf of T_V belongs to $\pi_V(A)$ and $|A| = k$. So the geodesic from $\beta_V(x)$ to $\beta_V(y)$ has a sub-segment of length at least $50Mk$ avoiding the points of valence more than 2. This subsegment contains a point α that necessarily separates $\beta_V(x)$ from $\beta_V(y)$ and either lies in $\{p_i^V\}$ or $\{r_V^W\}_{W \in \mathcal{V} \in \mathcal{V}}$, because such points form a $10Mk$ -net. The restriction to \mathcal{U}_1 is justified by the fact that for $W' \sqsubset W \sqsubset U$, we have that $\rho_V^{W'}$ coarsely coincides with ρ_V^W , so we can assume each r_V^W as above coincides with $r_V^{W'}$ for some $W' \in \mathcal{U}_1$ nested in V .

Choose any $x, y \in \mathcal{X}$ projecting M -close to b_V , and suppose by contradiction that $d_V(\beta_V(x), \beta_V(y)) > \tau$. 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 vertex of valence more than 2. 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 $\beta_W(x)$ and from $\beta_W(y)$.

Let \bar{T} be one of the two connected components of $T_W - \{p_i^W\}$. Then $\beta_W^{-1}(\bar{T})$ 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 . \square

Lemma 2.7. *There exists $\eta = \eta(M, k, \mathcal{X})$ such that the following holds. Let $p \in \mathcal{Y}$. Then the tuple $(b_V(p))$ is η -consistent.*

Proof. Let $U \pitchfork V$. If $U, V \in \mathfrak{S} - \mathcal{U}$, we are done because the corresponding coordinates b_U, b_V ($100Mk + E$)-coarsely coincide with those of, say, x_1 . If $U \in \mathcal{U}$ and $V \in \mathfrak{S} - \mathcal{U}$, then any point in T_V , whence also any point in $b_V(p)$, is $(100Mk + C + 2E)$ -close to ρ_V^U by Lemma 2.3 and the definition of \mathcal{U} , so we are done.

Now suppose that $U, V \in \mathcal{U}$. Let c_U be a point in T_U which is $10E$ -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 $100Mk$ -far from ρ_U^V and ρ_V^U respectively, 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 .

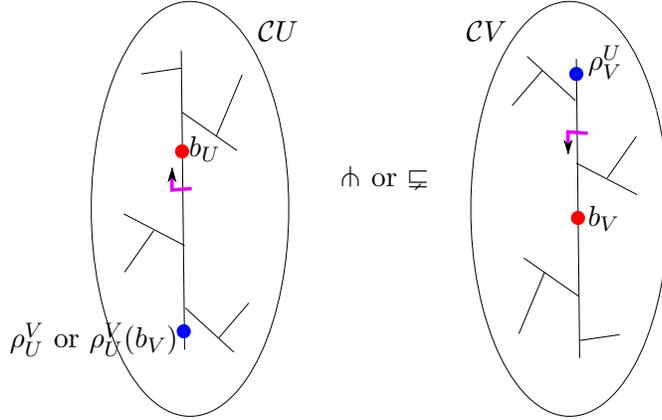


FIGURE 3. Proof of Lemma 2.7. $S_{W,i,U}(p), S_{W',i',V}(p)$ are shown as oriented halfspaces in the trees T_U, T_V .

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

100Mk–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 the bounded geodesic image axiom, $\rho_U^V(S_{W',i',V}(p))$ 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 . \square

2.3.2. Lemmas supporting the distance estimate. The next three lemmas support the distance estimate in the proof of Theorem 2.1.

The first lemma bounds projection distances from below in terms of the walls; it was used above to give a lower bound on $d_{\mathcal{X}}(\mathfrak{p}_A(p), \mathfrak{p}_A(q))$ in terms of the distance in the cube complex \mathcal{Y} between p and q .

Lemma 2.8. *Let $U \in \mathcal{U}$. For each $x, y \in H_\theta(A)$, we have $d_U(x, y) + 50EC\theta \geq |\{i : p_i^U \in [\beta_U(x), \beta_U(y)]\}|$. Moreover, if $\pi_U(x), \pi_U(y)$ are both C –close to $\pi_U(A)$, then $d_U(x, y) \geq |\{i : p_i^U \in [\beta_U(x), \beta_U(y)]\}|$.*

Proof. Let $x, y \in H_\theta(A)$. Recall that $\text{diam}(\pi_U(x) \cup \beta_U(x)) \leq 10(E + C + \theta)$, so $d_U(x, y) \geq d_U(\beta_U(x), \beta_U(y)) - 20(E + C + \theta)$. Hence $d_U(x, y) \geq d_{T_U}(\beta_U(x), \beta_U(y)) - 40EC\theta$. Therefore, $d_U(x, y) \geq |\{i : p_i^U \in [\beta_U(x), \beta_U(y)]\}| - 40EC\theta - 1$, as required. The “moreover” statement follows in a similar way using the fact that the p_i^U are M –far from leaves of T_U . \square

The next lemma is a simple application of Ramsey theory and the consistency property of an HHS. This lemma is used in tandem with the one above. It is also used below to control the number of separators associated to $p \in \mathcal{Y}$. Recall that p_i^U is said to be a separator for p if p_i^U separates $b_U(p)$ from $\beta_U(\mathfrak{p}_A)$ in the tree T_U .

Lemma 2.9. *There exists $N = N(\mathcal{X}) \geq 0$ so that for each $x \in H_\theta(A)$ there are at most N elements $U \in \mathcal{U}$ so that $d_U(\beta_U(x), \pi_U(A)) > 100E$.*

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 [BHS19, Lemma 2.1], c also bounds the maximum cardinality of a set of pairwise orthogonal elements. Given $x \in H_\theta(A)$, consider the set of $U \in \mathfrak{S}$ such that $d_U(x, A) > 100E$. 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 $U_1 \triangleleft U_2$ and $d_{U_l}(x, A) > 100E$ 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. \square

The next lemma bounds the number of separators in terms of M, k , and the constants ξ, τ . Since the proof is somewhat technical, we first give a heuristic discussion. We first show that if $p \in \mathcal{Y}$ has, say, T' separators, then there are at least $T'M/\xi$ elements $U \in \mathcal{U}$ that support separators (this is achieved by bounding the maximal number of separators supported on any given $U \in \mathcal{U}$). Lemma 2.9 shows that, for “most” such U , the point $\mathfrak{p}_A(p)$ projects in T_U close to some $\pi_U(x_j)$. So, if T' is too large, there is a specific pair x_j, x_k such that, in many U as above, $\mathfrak{p}_A(p)$ projects close to $\pi_U(x_j)$ and far from x_k . Lemma 1.6 then provides U_1, U_2 with these properties, both nested into some $V \in \mathfrak{S}$, such that $\rho_V^{U_1}, \rho_V^{U_2}$ are very far in \mathcal{CV} (in terms of M, ξ, τ). Applications of bounded geodesic image, consistency, and coherence of the orientation of walls corresponding to the 0–cube p allow us to conclude that $d_V(\pi_V(\mathfrak{p}_A(p)), b_V) > \xi$, which contradicts how the point $\mathfrak{p}_A(p)$ was chosen, namely as a realization point with constant ξ . For the last part of the argument, the reader will find it helpful to consult Figure 4.

Lemma 2.10. *There exists $T = T(M, k, \xi, \tau, \mathcal{X}, \mathfrak{S})$ such that for any $p \in \mathcal{Y}$ there exist at most T separators for p .*

Proof. We fix $p \in \mathcal{Y}$, after which, we can simplify our notation by writing b_V to mean $b_V(p)$. Recall from Lemma 2.7 that (b_V) is an η -consistent tuple, where η depends on M, k , and the global HHS constants, but is independent of p . Recall that the realization point $\mathfrak{p}_A(p) \in \mathcal{X}$ provided by Theorem 1.2 is characterized by the property that $\mathbf{d}_V(b_V, \mathfrak{p}_A(p)) \leq \xi$ for all $V \in \mathfrak{S}$, where ξ depends on η and the global HHS constants, but is independent of p .

First, for each $V \in \mathcal{U}$, we bound the number of separators p_i^V . Since each p_i^V separates b_V from $\beta_V(\mathfrak{p}_A(p))$, and the set of p_i^V in T_V is M -separated by construction, there are at most ξ/M separators with support V . Let Sep be the set of $V \in \mathcal{U}$ that support a separator for p .

By the previous paragraph, it remains to bound the cardinality $|\text{Sep}|$ of Sep . Now, Lemma 2.9 provides a uniform constant N so that there are at most N elements of \mathcal{U} where $\mathfrak{p}_A(p)$ projects $100E$ -far from every element of A .

Suppose that $|\text{Sep}| > N + N_0k(k-1)$, where $N_0 = N_0(M, \xi, \mathcal{X}, \mathfrak{S})$ will be chosen momentarily. (If the preceding inequality does not hold, then we have bounded $|\text{Sep}|$ independently of p , as required.) This lower bound implies that there are at least $N_0k(k-1)$ elements $V \in \text{Sep}$ where $\beta_V(\mathfrak{p}_A(p))$ is $100E$ -close to $\pi_V(A)$. Hence, there exists $x_j \in A$ so that there are at least $N_0(k-1)$ elements $V \in \text{Sep}$ where $\beta_V(\mathfrak{p}_A(p))$ is $100E$ -close to $\pi_V(x_j)$.

Now, for each such V , we have a separator p_i^V separating $\beta_V(\mathfrak{p}_A(p))$ from b_V , and necessarily lying M -far from $\pi_V(x_j)$, because of how our net in T_V was chosen. Hence b_V separates $\pi_V(x_j)$ from some $\pi_V(x_\ell)$. Hence there is a pair x_j, x_ℓ and at least N_0 elements $U \in \mathcal{U}$ such that:

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

Now, let $L = 1000(M + \xi + \tau)$. Suppose we chose $N_0 = N_0(L)$, the constant from Lemma 1.6. So, if $|\text{Sep}| > N_0k(k-1) + N$, then Lemma 1.6 provides some $V \in \mathfrak{S}$ and two elements $U_1, U_2 \in \text{Sep}$ with the above two listed properties, such that $U_1 \subsetneq V$, and $U_2 \subsetneq V$, and $\mathbf{d}_V(r_V^{U_1}, r_V^{U_2}) > 10E + \xi$.

For $t = 1, 2$, there exists $p_{i_t}^{U_t}$ separating $\beta_{U_t}(\mathfrak{p}_A(p))$ (which is $100E$ -close to $\pi_{U_t}(x_j)$) from $\pi_{U_t}(x_\ell)$, so $\mathbf{d}_{U_t}(x_j, x_\ell) > M$.

By bounded geodesic image, the geodesic in T_V from $\beta_V(x_j)$ to $\beta_V(x_k)$ must pass E -close to $r_V^{U_1}$ and $r_V^{U_2}$. (So V is necessarily in \mathcal{U} .)

For concreteness, suppose U_1, U_2 are labeled so that $\mathcal{N}_E(\rho_V^{U_1})$ separates $\pi_V(x_j)$ from $\rho_V^{U_2}$ and $\pi_V(x_\ell)$.

Bounded geodesic image and consistency imply that $\beta_V(\mathfrak{p}_A(p))$ lies E -close to the connected component of $T_V - \mathcal{N}_E(r_V^{U_t})$ containing $\pi_V(x_j)$ for $t = 1, 2$. Indeed, this holds because $\beta_{U_t}(\mathfrak{p}_A(p))$ is $(M - 100E)$ -far in T_{U_t} from $\pi_{U_t}(x_\ell)$.

We now analyze two cases, according to how close b_V lies to the component Π of $T_V - \mathcal{N}_E(\rho_V^{U_2})$ containing $\pi_V(x_k)$.

Recall that b_V is the intersection of various subtrees $S_{W,s,V}$, each of which coarsely coincides with the projection to \mathcal{CV} of a halfspace belonging to the coherent orientation p .

The first case is where every such subtree lies E -close to Π . In this case, the intersection of the subtrees — which is by definition b_V — lies $10E$ -close to Π . Hence, by Lemma 2.6, b_V is contained in the $(10E + \tau)$ -neighborhood of Π . This implies that $\mathbf{d}_V(b_V, \rho_V^{U_1}) \geq L - (10E + \tau)$. Since we saw that $\beta_V(\mathfrak{p}_A(p))$ is E -close to the component of $T_V - \mathcal{N}_E(r_V^{U_1})$ containing $\pi_V(x_j)$, we have $\mathbf{d}_V(\mathfrak{p}_A(p), b_V) > L - (10E + \tau) - E$. By our choice of L , this quantity exceeds ξ , which contradicts the definition of $\mathfrak{p}_A(p)$.

So we must be in the second case, where some halfspace H belonging to the coherent orientation p projects to a tree S in T_V — necessarily containing b_V — that is E -far from Π . Hence, S and $\pi_V(x_j)$ lie in the same component of $T_V - \mathcal{N}_E(\rho_V^{U_2})$. So, by the consistency

and bounded geodesic image axioms, $\pi_{U_2}(H)$ is contained in the E -neighborhood in T_{U_2} of $\pi_{U_2}(x_j)$. But since $p_{i_2}^{U_2}$ is M -far from $\pi_{U_2}(x_j)$ and separates $\beta_{U_2}(\mathfrak{p}_A(p))$ (which is $100E$ -close to $\pi_{U_2}(x_j)$) from b_{U_2} , we see that $p_{i_2}^{U_2}$ separates $\pi_{U_2}(H)$ from b_{U_2} . But the coherence of the orientation p and the definition of b_{U_2} requires b_{U_2} to be contained in $\pi_{U_2}(H)$, which gives a contradiction.

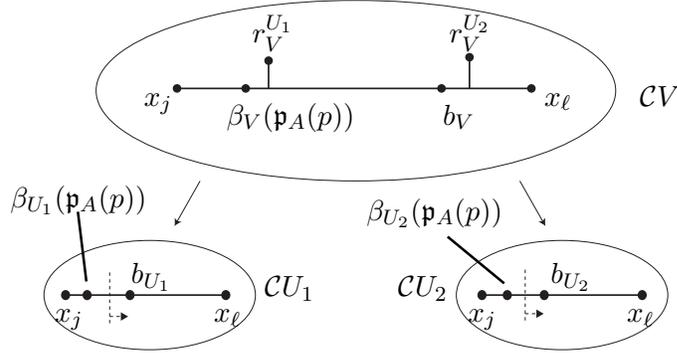


FIGURE 4. The proof of Lemma 2.10

We conclude that $|\text{Sep}| \leq N_0(L)k(k-1) + N$, which is independent of p . Hence there are at most $\xi(N_0(L)k(k-1) + N)/M$ separators for p , which is again independent of p because the realization constant ξ depends on p only to the extent that it depends on the consistency constant η for $(b_V(p))$, which was shown in Lemma 2.7 to be independent of p . \square

2.3.3. Walls cross if and only if orthogonal. We now check that the walls $\mathcal{L}_i^U, \mathcal{L}_j^V$ cross if and only if $U \perp V$. One direction, done in the first lemma, is essentially just the partial realization axiom for HHS. The other direction, which is the second lemma, relies on our specific choice of walls.

Lemma 2.11. *Suppose $U, V \in \mathcal{U}$ and $U \perp V$, and fix any $p \in \text{hull}_{\mathcal{CU}}(A)$, $q \in \text{hull}_{\mathcal{CV}}(A)$. Then there exists $x \in H_\theta(A)$ that coarsely projects to p in \mathcal{CU} and to q in \mathcal{CV} .*

Proof. By partial realization (Definition 1.1.(8) in [BHS19]), there exists $x' \in \mathcal{X}$ projecting E -close to p in \mathcal{CU} and q in \mathcal{CV} . 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. \square

Lemma 2.12 (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 2.11.

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 the image of A in each of the corresponding trees T_U, T_V 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, suppose $U \sqsubset V$. By construction, p_j^V is M -far from ρ_V^U , so we can choose a halfspace H associated to \mathcal{L}_j^V such that $\pi_V(H)$ contains a point $\pi_V(x_\ell)$ and is disjoint from $\mathcal{N}_E(\rho_V^U)$. Consistency and bounded geodesic image imply that $\pi_U(H)$ is E -close to $\pi_U(x_\ell)$ and hence M -far from p_i^U . Thus we can choose a halfspace H' for \mathcal{L}_i^U such that $\pi_U(H) \cap \pi_U(H') = \emptyset$, so $H \cap H' = \emptyset$, and hence \mathcal{L}_i^U and \mathcal{L}_j^V do not cross. \square

2.3.4. *The map \mathfrak{p}_A is quasimedial.* The next lemma is used to show that \mathfrak{p}_A is quasimedial, i.e., that it takes medians in the cube complex \mathcal{Y} (whose 1-skeleton is necessarily a median graph) close to coarse medians in \mathcal{X} .

Lemma 2.13. *There exists $C_1''' = C_1'''(\mathcal{X}, k, M)$ so that \mathfrak{p}_A is C_1''' -quasimedial.*

Proof. Let $\mu: \mathcal{X}^3 \rightarrow \mathcal{X}$ be the coarse median map. Recall from [BHS19] that μ is characterized by the following property: for all $a, b, c \in \mathcal{X}$ and all $U \in \mathfrak{S}$, $\pi_U(\mu(a, b, c))$ uniformly coarsely coincides with a coarse median point in $\mathcal{C}U$ for $\pi_U(a), \pi_U(b), \pi_U(c)$.

Let $x, y, z \in \mathcal{Y}$, and let m be their median. By Remark 1.32, 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 that 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 1.18. Hence $\mu(\mathfrak{p}_A(x), \mathfrak{p}_A(y), \mathfrak{p}_A(z))$ uniformly coarsely coincides with $\mathfrak{p}_A(q)$ for some $q \in \mathcal{Y}$.

Hence there exists $q \in \mathcal{Y}$ such that

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

and establishes that this distance 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. \square

At this point, we have proved all of the lemmas supporting Theorem 2.1.

2.4. Application to coarse median rank and hyperbolicity. In [BHS19, 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.

The following corollary bounds the median space rank of any asymptotic cone of \mathcal{X} ; see Proposition 1.35.

Corollary 2.14. *Suppose that \mathcal{X} is asymphoric. Then any $CAT(0)$ cube complex \mathcal{Y} from Theorem 2.1 satisfies $\dim \mathcal{Y} \leq \nu$, 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 $\mathcal{Y} \rightarrow H_\theta(A)$ be the C_1 -quasimedial (C_1, C_1) -quasi-isometry provided by Theorem 2.1. By Lemma 2.12, $\dim \mathcal{Y}$ 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 $\geq M$, such subsets have cardinality bounded by ν . This proves Corollary 2.14.

Moreover, $\mathcal{Y}^{(0)} \rightarrow H_\theta(A)$ is a quasimedial 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, $\dim \mathcal{Y} \leq \nu$. Hence \mathcal{X} is coarse median of rank ν . \square

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 exist D and δ so that

- any pair of points of X is joined by a (D, D) -quasi-geodesic, and
- (D, D) -quasi-geodesic triangles are δ -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 HHSs which are hyperbolic, which we use below in the proof of Lemma 4.6:

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

- \mathcal{X} is coarse median of rank ≤ 1 , and is thus hyperbolic;
- (Bounded orthogonality) There exists $q \in \mathbb{R}$ so that $\min\{\text{diam}(CU), \text{diam}(CV)\} \leq q$ 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]. \square

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 CAT(0) cube complex, i.e., a tree.

3. QUASIFLATS AND ASYMPTOTIC CONES

Fix an asymphoric hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ of rank ν and let \mathcal{X} be an asymptotic cone of \mathcal{X} . According to Proposition 1.35, the coarse median map on \mathcal{X} limits to a median map on \mathcal{X} making it into a topological median space of rank at most ν . By the same proposition, after changing the metric on \mathcal{X} within its bilipschitz equivalence class, we can assume that \mathcal{X} , with its given median, is a median metric space.

With this setup in mind, we now outline this section. First, the goal is to show that given a quasiflat in \mathcal{X} , there are arbitrarily large balls contained in a uniform neighborhood of the hull of boundedly many points; this is made precise in Corollary 3.9, and this is what will allow us to apply Huang’s quasiflat theorem for CAT(0) cube complexes [Hua14b] to describe quasiflats in HHSs. Subsection 3.1 contains preliminary lemmas that relate ultralimits of objects in \mathcal{X} defined in terms of the HHS structure to objects in \mathcal{X} defined in terms of the median structure.

The content of Lemma 3.3 and Proposition 3.4 is best explained in reversed order: In Proposition 3.4, we argue that there are balls of large radius R in quasiflats in \mathcal{X} that stay ϵR -close to hulls of finitely many points, for a fixed small $\epsilon > 0$. Taking ultralimits, this gives a bilipschitz flat in an asymptotic cone that stays within bounded distance of a certain median convex subspace, and Lemma 3.3 says that this means that, in fact, the flat is contained in the convex subspace. This corresponds to an improvement from “ ϵR ” to “ $o(R)$ ”.

We then need to further improve this to “ $O(1)$ ”, which is performed in Proposition 3.5 by shrinking the previously found balls. Further discussion of the various statements and the corresponding proofs can be found below.

3.1. Ultralimits of hulls and some median preliminaries. Given m, m' in a median space M , we let $\text{hull}(m, m')$ denote the set of $z \in M$ for which the median of m, m', z is z . (Note that $\text{hull}(m, m') = [m.m']$, where $[m.m']$ is the *median interval*, defined just as before. The term “median interval” is more standard, but we think of median intervals as convex hulls of pairs of points, which explains our choice of notation.)

Fix a hierarchically quasiconvex subspace $A \subseteq \mathcal{X}$ and points $p, q \in A, x \in \mathcal{X}$. Note that the coarse median of (p, q, x) lies uniformly close to A (see e.g., [BHS19, Section 7] or [RST18, Section 5]) — this easily yields the first assertion of the following lemma, which we use freely throughout this section.

Lemma 3.1. *For any κ , the ultralimit of any sequence of κ -hierarchically quasiconvex subspaces is median convex. Moreover, if (A_n) is a sequence of κ -hierarchically quasiconvex subspaces and $\mathbf{A} \subset \mathcal{X}$ is their ultralimit, then the maps $\mathfrak{g}_{A_n} : \mathcal{X} \rightarrow A_n$ limit to the median gate map $\mathfrak{g} : \mathcal{X} \rightarrow \mathbf{A}$.*

Proof. We prove the assertion about gates, as the other facts are already established, as noted above. Fix $\mathbf{x} \in \mathcal{X}$, represented by a sequence (x_n) in \mathcal{X} . Fix $\mathbf{a} \in \mathbf{A}$, represented by a sequence (a_n) . For each n , let $b_n = \mathfrak{g}_{A_n}(x_n)$, and let \mathbf{b} be represented by (b_n) .

By the definition of the gate and the coarse median, the coarse median of a_n, b_n, x_n is uniformly close to b_n , so the median of $\mathbf{a}, \mathbf{b}, \mathbf{x}$ is \mathbf{b} . Hence the median interval between \mathbf{x} and any point in \mathbf{A} contains \mathbf{b} ; it follows immediately from the definition of gate that $\mathbf{b} = \mathfrak{g}(\mathbf{x})$. \square

We will also tacitly use the next lemma throughout this section. It states that the (median) convex hull of a pair of points in an asymptotic cone of \mathcal{X} arises as a limit of θ -hulls of pairs of points in \mathcal{X} .

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

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

$$\lim_{\omega} H_{\theta}(x_n, y_n) \subseteq \text{hull}(\mathbf{x}, \mathbf{y}).$$

To prove the other containment, suppose $\mathbf{z}' \in \text{hull}(\mathbf{x}, \mathbf{y})$ (and whence, by definition of the hull, that $\mathbf{z}' = m(\mathbf{x}, \mathbf{y}, \mathbf{z}')$), and let (z'_n) be a representative for \mathbf{z}' . Let $z_n = m(x_n, y_n, z'_n) \in H_{\theta}(x_n, y_n)$ and note that this implies $\mathbf{z} = m(\mathbf{x}, \mathbf{y}, \mathbf{z}')$, where \mathbf{z} is the point represented by (z'_n) . Since \mathcal{X} is a median space, the median of a triple is unique and thus $\mathbf{z}' = \mathbf{z}$; whence $\mathbf{z}' \in \lim_{\omega} H_{\theta}(x_n, y_n)$. \square

3.2. Bilipschitz flats in asymptotic cones. The next lemma relies on results of Bowditch about *culubated* subsets of median metric spaces [Bow18b]. The import of the lemma is the following. Consider a top-rank bilipschitz flat \mathbf{F} in \mathcal{X} and a median-convex subspace \mathbf{H} arising as an ultralimit of hierarchically quasiconvex subspaces of \mathcal{X} . If \mathbf{F} lies in a uniform neighborhood of \mathbf{H} , then it must actually be contained in \mathbf{H} . This will be applied in the proof of Proposition 3.5 in the case where \mathbf{H} is a limit of θ -hulls of finite sets in \mathcal{X} of bounded cardinality.

Roughly, the idea of proof is as follows. If the bilipschitz flat \mathbf{F} was median convex, we would have gate maps on both \mathbf{F} and \mathbf{H} , and \mathbf{F} could only stay close to \mathbf{H} around $\mathfrak{g}_{\mathbf{F}}(\mathbf{H})$, which then needs to be the whole of \mathbf{F} . Since \mathbf{F} is top-dimensional, and $\mathfrak{g}_{\mathbf{F}}(\mathbf{H})$ is one of the factors of a product subspace of \mathcal{X} (in view of Lemma 1.20), the other factor has to be trivial.

On the other hand, the other factor being trivial is the same as \mathbf{F} being contained in \mathbf{H} . Now, \mathbf{F} need not be median convex, and to deal with this we rely on results from [Bow18b] that, roughly, give us a decomposition of (large portions of) the quasiflat into “blocks,” each of which is median convex, and we then consider chains of such blocks.

Lemma 3.3 (Close to convex implies contained in convex). *Let \mathcal{X} be an asymptotic cone of \mathcal{X} and let $\mathbf{F} \subseteq \mathcal{X}$ be a bilipschitz ν -flat. Let \mathbf{H} be an ultralimit of uniformly hierarchically quasiconvex subsets of \mathcal{X} and suppose that \mathbf{F} is contained in a neighborhood of \mathbf{H} of finite radius. Then $\mathbf{F} \subseteq \mathbf{H}$.*

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

By [Bow18b, Proposition 1.2], \mathbf{F} is cubulated in the sense of [Bow18b], which means that there are arbitrarily large balls in \mathbf{F} each of which is contained in a finite union of blocks. By [Bow18b, Proposition 3.3], this implies that there are arbitrarily large balls B in \mathbf{F} with the following property: B is contained in a subset of \mathbf{F} which is a union of blocks whose pairwise intersections are each either empty or a common face. We let \mathbf{F}' be such a union of blocks which contains a ball around $p \in \mathbf{F}$ of radius r much larger than $\sup_{x \in \mathbf{F}} d(x, \mathbf{H})$.

After possibly subdividing the cubulation of \mathbf{F}' , there is a ν -block B_0 of \mathbf{F}' containing p and disjoint from \mathbf{H} . After subdividing, we can assume that each side of B_0 has length bounded by some ℓ much smaller than r .

Being a block, B_0 is the median interval between a pair of opposite corners of B_0 . So, by Lemma 3.2, B_0 is the ultralimit \mathbf{H}_0 of a sequence $(H_\theta(c_n, d_n))$ of θ -hulls of pairs of points.

As noted in Definition 1.36, $\mathfrak{g}_{\mathbf{H}_0}(\mathbf{H})$ is a median convex subspace. So, $\mathfrak{g}_{\mathbf{H}_0}(\mathbf{H})$ is a sub-block B' of B_0 .

On the other hand, by hypothesis, \mathbf{H} is the limit of uniformly hierarchically quasiconvex subspaces H_n of \mathcal{X} . By Lemma 3.2, $\mathfrak{g}_{\mathbf{H}_0}(\mathbf{H})$ is the limit of the subspaces $\mathfrak{g}_{H_\theta(c_n, d_n)}(H_n)$. By Lemma 1.20.(3), $\mathfrak{g}_{H_\theta(c_n, d_n)}(H_n)$ is coarsely contained in a quasi-isometrically embedded copy of $\mathfrak{g}_{H_\theta(c_n, d_n)}(H_n) \times I_n$, where I_n is a θ -hull. So, taking limits, we see that, if B' has dimension i , then there is an $(i + 1)$ -dimensional topologically embedded copy of $[0, 1]^{i+1}$ in \mathcal{X} . This implies $i < \nu$.

For any codimension-1 face B_2 of B_0 not intersecting B' , there exists a block B'_1 whose intersection with B_0 is B_2 . So, $B_1 = B_0 \cup B'_1$ is a block by [Bow18b, Lemma 3.2]. We claim $\mathfrak{g}_{B_1}(\mathbf{H}) = \mathfrak{g}_{B_0}(\mathbf{H})$, which implies that B_1 is also disjoint from \mathbf{H} .

To prove the claim, note that $B' = \mathfrak{g}_{B_0}(\mathbf{H}) = \mathfrak{g}_{B_0}(\mathfrak{g}_{B_1}(\mathbf{H}))$. (It is a general fact about median metric spaces, following directly from the definition of a gate, that if A, C are median-convex closed subspaces and $A \subset C$, then $\mathfrak{g}_A = \mathfrak{g}_A \circ \mathfrak{g}_C$.)

Now, since B_0 is a sub-block of the block B_1 , and B_0 intersects the closure of its complement in B_1 along a common codimension-1 face, and B_1 is median-isomorphic to a finite product of intervals with the ℓ_1 -metric (by the definition of a block), $\mathfrak{g}_{B_0}|_{B_1}$ is just the natural retraction. So, this map is one-to-one on B' , and the claim follows.

Now proceed inductively until we find a block B_m that we cannot extend to a block B_{m+1} using the procedure above, implying that we reached the boundary of \mathbf{F}' . By induction, $\mathfrak{g}_{B_m}(\mathbf{H}) = \mathfrak{g}_{B_0}(\mathbf{H})$.

Hence $\mathfrak{g}_{B_m}(\mathbf{H}) = B'$, since we had $B' = \mathfrak{g}_{B_0}(\mathbf{H})$. Let $q \in B_m$ lie in the boundary of \mathbf{F}' . Then $d_{\mathcal{X}}(q, B')$ is at least $d_{\mathcal{X}}(q, p) - \nu\ell$, which exceeds $\sup_{x \in \mathbf{F}} d(x, \mathbf{H})$. Hence there exists $h \in \mathbf{H}$ with $d_{\mathcal{X}}(h, q) < d_{\mathcal{X}}(B', h)$. This contradicts that $\mathfrak{g}_{B_m}(\mathbf{H}) = B'$. (Here we are using the median metric $d_{\mathcal{X}}$, for which the notions of median gate and closest-point projection coincide, by the definition of a median gate.) This is the required contradiction. \square

3.3. Quasiflats and hulls. As mentioned above, we now argue that given a quasiflat in \mathcal{X} there are balls of large radius R that stay ϵR -close to hulls of finitely many points, for a fixed

small $\epsilon > 0$. Once again we use [Bow18b, Proposition 1.2], which provides a subdivision of (large portions) of the ultralimit of the quasiflat into blocks, and then use the fact that each such block is the ultralimit of hulls of pairs of points.

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

Proof. The proof has two parts.

Choosing N : Let \mathcal{X} be a *fixed* asymptotic cone of \mathcal{X} with observation points a constant sequence $(F(0))$. Let $\mathbf{F}: \mathbb{R}^\nu \rightarrow \mathcal{X}$ be the corresponding ultralimit of F . Let \mathbf{B} be a ball of radius 1 in \mathbb{R}^ν . By [Bow18b, Proposition 1.2], $\mathbf{F}(\mathbf{B})$ 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, $\mathbf{F}(\mathbf{B})$ is contained in the ultralimit of hulls of pairs of points. Thus, $\mathbf{F}(\mathbf{B})$ 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).

Remark on non-uniformity of N : We remark that, for the purposes of this proof, N is allowed to depend on the particular quasiflat F , not just the quasi-isometry constants. We are also allowing N to depend on our choice \mathcal{X} of asymptotic cone. Bowditch's proposition (Proposition 1.2 in [Bow18b]) provides only that $\mathbf{F}(\mathbf{B})$ is contained in a *finite* union of blocks, but does not bound the number; for our result we only need finiteness.

Conclusion: 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 $B(0, R)$ of sufficiently large radius R , we have that $F(B(0, R))$ cannot be contained in $\mathcal{N}_{\epsilon R}(H_\theta(A))$ for any $A \subseteq \mathcal{X}$ with $|A| \leq N$. Let $B_n = B(0, R_n)$, where (R_n) is the scaling factor of the asymptotic cone \mathcal{X} fixed above. Then \mathbf{B} is the ultralimit of the B_n . The fact that $\mathbf{F}(\mathbf{B})$ 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 $\mathcal{N}_{\epsilon R_n}(H_\theta(A_n))$, a contradiction. \square

The following is the most technical proposition of this section, and it says that by shrinking the balls provided by Proposition 3.4, we obtain balls contained in a uniform neighborhood of hulls of boundedly many points. The rough reason for this is the following. In view of Lemma 3.3, in any asymptotic cone the ultralimit of the balls is contained in the ultralimit of the hulls; this means that the distance from the flat to the hulls grows more slowly than any superlinear function. From this we deduce the distance is bounded. To make this work, we must consider only asymptotic cones where the ultralimit of the balls is a bilipschitz flat, so the observation point must be deep in the balls; in the proof we deal with this by using balls of half the radius to ensure this holds in the relevant asymptotic cones.

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 $R \geq R_0$ in \mathbb{R}^ν , and let $F: B \rightarrow \mathcal{X}$ be a (K, K) -quasi-isometric embedding. Let $A \subseteq \mathcal{X}$ have $|A| \leq N$, 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}^ν , and (K, K) -quasi-isometric embeddings $F_m: B_m \rightarrow \mathcal{X}$,
- subsets $A_m \subseteq \mathcal{X}$ with $|A_m| \leq N$ and

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

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

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

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

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

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

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

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

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

Remark 3.7. The proof follows from Claim 3.6 by an application of the principle from nonstandard analysis called *underspill*, which says that if a predicate is true for all infinitesimal positive non-standard reals, then it is also true for all sufficiently small standard reals.

Since we do not wish to require familiarity with non-standard analysis, rather than invoking this principle we instead provide a self-contained argument in the language of ultrafilters. Since our argument is a translation of the non-standard analysis argument, rather than providing a convoluted heuristic explanation for how this argument works, we refer the reader who would like to do more than check that the argument is formally correct to Tao's excellent blog post [Tao], which explains all the relevant concepts. We note, though, that this argument is the usual one which is used to prove that ultrapowers are saturated models and also in proving the nonstandard formulation of continuity, see [Tao, Proposition 11], which is a typical application of underspill.

For each $n \in \mathbb{N}$, let \mathcal{U}_n be the set of $m \geq n$ for which there exists $t_{m,n}, \alpha_{m,n} \in \mathbb{R}_+$ so that $\ell_m(t_{m,n}) \geq n$ and $\alpha_{m,n} \geq n$ 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 $\ell(\lambda - \alpha\ell(\lambda)) \ll \ell(\lambda)$. 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 $\ell_m(R_m) \leq R_m/(4\alpha_0)$ (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 3.8. *Either $\ell(R_m^j) \leq \ell_m(R_m)/2^j$ or there exists $i \leq j$ with $\ell_m(R_m^i) < T_0$.*

Proof of Claim 3.8. We argue by induction on j . Suppose that R_m^j satisfy $\ell_m(R_m^j) \leq \ell_m(R_m)/2^j$ and $\ell_m(R_m^j) \geq T_0$. 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 $\alpha_m^j \geq T_0$. Hence, the claim gives $\ell_m(R_m^{j+1}) \leq \ell_m(R_m^j)/2 \leq \ell(R_m)/2^{j+1}$, as required. \square

In either of the two cases provided by Claim 3.8, there exists j with $\ell_m(R_m^j) < T_0$. This implies $\ell_m(R_m/2) < T_0$, and hence $\ell(\sigma/2) < T_0$, as required. \square

Combining Proposition 3.4 and Proposition 3.5, one gets:

Corollary 3.9. *For every quasi-isometric embedding $f: \mathbb{R}^n \rightarrow \mathcal{X}$, 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 $|A_R| \leq N$ 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 . The main goal of this section is to prove the quasiflats theorem, Theorem A, which says that quasiflats in \mathcal{X} are at bounded Hausdorff distance from a finite union of standard orthants. After some preliminary work on orthants, we complete the proof in Subsection 4.3. Then, we prove two further results giving more quantitative control on quasiflats in \mathcal{X} in terms of the number of orthants needed (Theorem 4.14) and the Hausdorff distance between the quasiflat and not quite the union of the orthants, but rather the hull of the union (Lemma 4.17).

4.1. Orthants in \mathcal{X} . We fix once and for all a constant D so that, for any $U \in \mathfrak{S}$, any two points in F_U are connected by a D -hierarchy path. (Such a constant is provided by Theorem 1.4.)

We now discuss *standard orthants* in \mathcal{X} , which are one of the basic objects in the statement of Theorem A.

Definition 4.1 (Standard orthant, standard flat, standard partial flat). Let U_1, \dots, U_k be pairwise orthogonal elements of \mathcal{X} . Recall that we have a quasimedial quasi-isometric embedding $F_{U_1} \times \dots \times F_{U_k} \rightarrow \mathcal{X}$, as described in Section 1.1.1, with constants independent of the U_i .

For each $i \leq k$ and each $x \in \prod_{j \neq i} F_{U_j}$, the image of $F_{U_i} \times \{x\}$ is a (uniformly) hierarchically quasiconvex subset which, abusing notation slightly, we also denote $F_{U_i} \times \{x\}$, or simply by F_{U_i} when the choice of parallel copy is not important.

For each i , let γ_i be a D -hierarchy ray in F_{U_i} with the property that $\pi_{U_i}(\gamma_i)$ is unbounded. We call the image of $\gamma_1 \times \dots \times \gamma_k \subseteq F_{U_1} \times \dots \times F_{U_k}$ under the standard embedding a *standard k -orthant* in \mathcal{X} with support set $\{U_i\}$.

A *standard orthant* is a standard ν -orthant, i.e., a standard k -orthant of maximum possible dimension.

Similarly, given U_1, \dots, U_k as above, suppose we have for each $i \leq k$ a path γ_i in F_{U_i} such that γ_i is either a D -hierarchy ray or a bi-infinite D -hierarchy path such that $\pi_{U_i}(\gamma_i)$ is unbounded. Then the image of $\gamma_1 \times \dots \times \gamma_k$ is a *standard k -partial flat*, or a *standard partial flat* if $k = \nu$. If every γ_i is bi-infinite, then we use the term *standard k -flat*, or *standard flat* if $k = \nu$.

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

The next lemma says that top-dimensional standard orthants in an asymphoric HHS are hierarchically quasiconvex (with uniform hierarchical quasiconvexity function). Here, an analogy to the CAT(0) cube complex situation is again instructive. If Π is a CAT(0) cube complex, and $O \subset \Pi$ is a cubical orthant, then although O is ℓ_1 -isometrically embedded (i.e., its 0-skeleton is a 1-connected median subalgebra) by definition, it need not be convex: picture the case where $\Pi = \mathbf{R}^2$ and O is the ray with 0-skeleton consisting of points

$\{(n, n), (n, n + 1) : n \in \mathbb{N}\}$. On the other hand, if O has the property that $\dim O = \dim \Pi$, then O cannot contain the “corner” of any cube of Π that is “missing” in O , i.e., O is convex. This cubical fact is important in Huang’s work [Hua14b]. The final assertion of the next lemma is analogous.

Lemma 4.3 (Top dimensional orthants are hierarchically quasiconvex). *Consider a standard k -orthant O whose support set $\{U_i\}$ has the property that, for some C , we have $\min\{\text{diam}_{\mathcal{CU}}(\pi_U(O)), \text{diam}_{\mathcal{CV}}(\pi_V(O))\} \leq C$ whenever $U, V \subseteq U_i$ are orthogonal and $i \leq k$. 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.

Remark 4.4. Lemma 4.3 holds when the standard orthant O is replaced by a standard flat or standard partial flat; the exact same proof works, except with some of the rays replaced by bi-infinite paths. The main lemma being used is Lemma 4.5, which is stated for arbitrary hierarchy paths.

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 $\{U_1, \dots, U_k\}$ 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{CU} , 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 $U \subseteq U_i$, this follows from hierarchical quasiconvexity of γ_j , i.e., Lemma 4.5. \square

The next lemma supports the preceding one. It gives a sufficient condition for a hierarchy path to be hierarchically quasiconvex. The reader familiar with the work of Huang may find it useful to compare this lemma with the notion of a “straight” geodesic in a CAT(0) cube complex, defined in [Hua14b].

Lemma 4.5 (“Straight” hierarchy paths). *Let $\gamma : I \rightarrow \mathcal{X}$ 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 $M \geq \max\{C, M_0\}$, 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 -quasimedial (C_1, C_1) -quasi-isometric embedding $\mathcal{C}(x, y) \rightarrow \mathcal{X}$ 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 $M \geq C$ imply that $\dim \mathcal{C}(x, y) \leq 1$. 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 $i \geq 0$ 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. \square

In the proof of Theorem A, we will construct a quasimedial quasi-isometric embedding of a CAT(0) cube complex into \mathcal{X} . Huang's theorem will provide cubical orthants in the CAT(0) cube complex, so we need to prove that the image of each cubical orthant is coarsely a standard orthant. For that, we will use the following lemma:

Lemma 4.6. *Let O be an ν -dimensional cubical orthant with a quasimedial quasi-isometric embedding $q: O \rightarrow \mathcal{X}$. Then there is a standard orthant $Q \subset \mathcal{X}$ with $d_{\text{haus}}(q(O), Q) < \infty$.*

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

Related points and pairs: We say that $x, y \in O$ are i -related, for $1 \leq i \leq \nu$, 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 $1 \leq i \leq \nu$, let \mathcal{U}_i be the collection of all $U \in \mathfrak{S}$ so that there exist i -related $x, y \in O$ with $d_U(q(x), q(y)) \geq M$. For any K , we also let $\mathbf{Rel}_K(q(O)) = \{U \in \mathfrak{S} : \text{diam}_{\mathcal{C}U}(\pi_U(q(O))) \geq K\}$.

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

Claim 4.7. *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*

- $d_U(x, y) \leq C$ and $d_U(x', y') \leq C$, or
- $d_U(x, x') \leq C$ and $d_U(y, y') \leq C$.

Proof of Claim 4.7. Let $m: O^3 \rightarrow O$ 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 4.8. *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 4.8. Consider distinct i, j , an i -related pair x, y and some U with $d_U(q(x), q(y)) \geq M$, and a j -related pair w, z and some V so that $d_V(q(w), q(z)) \geq M$.

Provided $M \geq 10(\nu - 1)C$, applying Claim 4.7 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:

- $d_V(q(x), q(x')) \geq M/2$ and $d_V(q(y), q(y')) \geq M/2$;
- $d_U(q(x), q(y)) \geq M$ and $d_U(q(x'), q(y')) \geq M/2$.

Claim 4.7 implies that $d_U(q(x), q(x')) \leq C$, $d_U(q(y), q(y')) \leq C$ and $d_V(q(x), q(y)) \leq C$, $d_V(q(x'), q(y')) \leq C$.

For M large enough, this implies that $U \perp V$. Indeed, if $U = V$, then the triangle inequality yields $4C \geq M/2$, 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_V^U , contradicting consistency. A similar contradiction arises if U, V are \sqsubset -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 quasimedial 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 Lemma 1.37. By [DHS17, Lemma 3.3], there exist $U_1^i, \dots, U_{k_i}^i$ so that

$\pi_{U_j^i}(\gamma'_i)$ is unbounded. Moreover, by the same lemma, for $1 \leq i \leq \nu$, $1 \leq j < j' \leq k_i$, we have $U_j^i \perp U_{j'}^i$.

Since each $U_j^i \in \mathcal{U}_i$, Claim 4.8 and the fact that \mathcal{X} has rank ν implies that $k_i = 1$ for each i . To streamline notation, let $U_i = U_1^i$.

Since $\{U_1, \dots, U_\nu\}$ is a pairwise-orthogonal set, the following holds for all $i \leq \nu$: if $U, V \not\subseteq U_i$ have $\text{diam}(\mathcal{C}U), \text{diam}(\mathcal{C}V) > E$, then $U \not\perp V$, for otherwise $\{U_1, \dots, U_{i-1}, U, V, U_{i+1}, \dots, U_\nu\}$ 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 $d_{\mathcal{C}U}(q(p), q(p_i)) \leq C'$ 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 $j \neq i$.

Now let $p \in O$. 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. \square

4.2. Coarse intersections of orthants. In this subsection we study coarse intersections of orthants. This is mostly needed for the next section, but we need Lemma 4.11 in the proof of Theorem 4.14.

Definition 4.9 (Coarse intersection). Let $A, B \subset \mathcal{X}$. Suppose that there exists R_0 so that for any $R, R' \geq R_0$, we have $d_{\text{haus}}(\mathcal{N}_R(A) \cap \mathcal{N}_R(B), \mathcal{N}_{R'}(A) \cap \mathcal{N}_{R'}(B)) < \infty$. 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 $A \tilde{\cap} B$.

In the next lemma, we show that, for pairs of hierarchically quasiconvex sets, an R_0 as in the definition above exists, and so the coarse intersection is well-defined. This is one of the places where we use the bridge lemma (Lemma 1.20).

Lemma 4.10 (Coarse intersections coarsely coincide with gates). *For all κ, r , there exists R_0 such that the following holds. Let A, B be κ -hierarchically quasiconvex and suppose $d(A, B) \leq r$. Then for all $R, R' \geq R_0$, we have $d_{\text{haus}}(\mathcal{N}_R(A) \cap \mathcal{N}_R(B), \mathcal{N}_{R'}(A) \cap \mathcal{N}_{R'}(B)) < \infty$, so $A \tilde{\cap} B$ is well-defined. Moreover, there exists $K = K(\kappa, r)$ such that $A \tilde{\cap} B$ is at Hausdorff distance at most K from $\mathfrak{g}_A(B)$.*

Proof. By Lemma 1.20.(3), there exists K_1 , depending only on $\kappa(0)$ and E , and a (K_1, K_1) -quasi-isometric embedding $f: \mathfrak{g}_A(B) \times H_\theta(\{a, b\}) \rightarrow \mathcal{X}$ such that $\text{im } f$ K_1 -coarsely coincides with $H_\theta(\mathfrak{g}_A(B) \cup \mathfrak{g}_B(A))$, where $a = \mathfrak{g}_A(b) \in \mathfrak{g}_A(B)$ and $b = \mathfrak{g}_B(a)$. By Lemma 1.27 (applied exchanging the roles of A and B), $d(a, b)$ is bounded in terms of κ and r . Hence, there exists R_1 , depending on κ, K_1 , and r , such that any point in $\mathfrak{g}_A(B) \subseteq A$ lies at distance at most R_1 from $\mathfrak{g}_B(A) \subseteq B$, and hence at distance at most R_1 from B . So, $\mathfrak{g}_A(B) \subset \mathcal{N}_{R_1}(A) \cap \mathcal{N}_{R_1}(B)$.

On the other hand, if $p \in \mathcal{N}_R(A) \cap \mathcal{N}_R(B)$ for some R , then apply Lemma 1.27 to find $K = K(\kappa)$ such that $d(p, A) \asymp_{K,K} d(p, \mathfrak{g}_A(p))$ and $d(A, B) \asymp_{K,K} d(\mathfrak{g}_A(p), \mathfrak{g}_B(\mathfrak{g}_A(p)))$. So, $d(p, \mathfrak{g}_B(\mathfrak{g}_A(p))) \leq_{K,K} R + r$. In other words, $d(p, \mathfrak{g}_B(A))$ is uniformly bounded (in terms of κ, R and r) and $d(p, \mathfrak{g}_A(B))$ is bounded similarly. So $\mathcal{N}_R(A) \cap \mathcal{N}_R(B)$ uniformly coarsely

coincides with $\mathfrak{g}_A(B)$, proving the second claim. Since any two neighborhoods of $\mathfrak{g}_A(B)$ coarsely coincide, the first claim follows. \square

The following lemma describes coarse intersections of orthants, which, as one might hope, turn out to be sub-orthants.

Lemma 4.11 (Coarse intersections of orthants). *Let O, O' be standard orthants in \mathcal{X} with supports $\{U_i\}_{i \leq \nu}, \{U'_i\}_{i \leq \nu}$. 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 $\{U_i\}_{i \leq \nu} \cap \{U'_i\}_{i \leq \nu}$.*

Proof. By Lemma 4.10, 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 γ'_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 1.20.(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{CU} 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 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 . \square

In the proof of Theorem 5.7 below, we will need the following version of the above lemma, stated for coarse intersections of standard flats instead of standard orthants.

Lemma 4.12 (Coarse intersection of standard flats). *Let F, F' be standard flats in \mathcal{X} with supports $\{U_i\}_{i=1}^\nu$ and $\{U'_i\}_{i=1}^\nu$ respectively. Then $F \tilde{\cap} F'$ is well-defined, and coarsely coincides with $\mathfrak{g}_F(F')$. Moreover, suppose that $\{U_i\} \cap \{U'_i\} = \{U\}$ for some $U \in \mathfrak{S}$. Then $F \tilde{\cap} F'$ is either a bounded set or coarsely coincides with a standard 1-orthant or standard 1-flat with support $\{U\}$.*

Similarly, if $\{U_i\} \cap \{U'_i\} = \{U, V\}$ for some (necessarily orthogonal) $U, V \in \mathfrak{S}$, then $F \tilde{\cap} F'$ coarsely decomposes as the product of two hierarchically quasiconvex subspaces α, β , each of which is either bounded or coarsely coincides with a standard 1-orthant or standard 1-flat.

Proof. The standard flats F, F' are uniformly hierarchically quasiconvex by Remark 4.4. Lemma 4.10 implies that $F \tilde{\cap} F'$ is well-defined and coarsely coincides with $\mathfrak{g}_F(F')$. So, we just need to show that $\mathfrak{g}_F(F')$ is a standard 1-orthant or 1-flat with support $\{U\}$, or $\mathfrak{g}_F(F')$ is bounded. For each $i \leq \nu$, let γ_i be the hierarchy path in F_{U_i} which is the i^{th} factor of F , and define γ'_i analogously for F' . Re-labeling if necessary, let $U = U_1 = U'_1$. Note that by Lemma 4.5, each γ_i, γ'_i is uniformly hierarchically quasiconvex. Indeed, $\pi_V(\gamma_i)$ has uniformly bounded diameter unless $V \subseteq U_i$. But if $V, W \subseteq U_i$ are orthogonal, then $\{V, W\} \cup \{U_j\}_{j \neq i}$ is a pairwise-orthogonal set of $\nu + 1$ elements, so by asymphoricity, \mathcal{CV} (say) has diameter at most E , so the same is true of $\pi_V(\gamma_i)$. Hence Lemma 4.5 applies.

Let $\alpha = \mathfrak{g}_{\gamma_1}(\gamma'_1)$. Arguing as in the proof of Lemma 4.11 shows that α , which coarsely coincides with $F \tilde{\cap} F'$, is either bounded or coarsely coincides with a 1-orthant or 1-flat. This proves the first assertion.

The second assertion follows similarly. Again, $F \tilde{\cap} F'$ coarsely coincides with $\mathfrak{g}_F(F')$ by Lemma 4.10, so it suffices to show that $\mathfrak{g}_F(F')$ coarsely coincides with a product $\alpha \times \beta$ as in the statement. Label the supports of F, F' so that $U = U_1 = U'_1$ and $V = U_2 = U'_2$. Let

$\alpha = \mathfrak{g}_{\gamma_1}(\gamma'_1)$ and let $\beta = \mathfrak{g}_{\gamma_2}(\gamma'_2)$. Then argue as in Lemma 4.11 to see that $\alpha \times \beta$ coarsely coincides with $\mathfrak{g}_F(F')$. \square

4.3. Quasiflats theorem. We are now ready to prove Theorem A, which we restate as:

Theorem 4.13. *Let \mathcal{X} be an asymphoric HHS of rank ν and let $f: \mathbb{R}^\nu \rightarrow \mathcal{X}$ be a quasi-isometric embedding. Then there exists a finite set of standard orthants $Q_i \subseteq \mathcal{X}$ for $1 \leq i \leq k$, for which:*

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

Proof. Let L, N be as in Corollary 3.9. Then there exists an increasing unbounded sequence $R_1 < R_2 < \dots$ and sets $A_i \subseteq \mathcal{X}$ of cardinality at most N for which the following holds. Let B_i be the ball in \mathbb{R}^ν 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 $c_i: \mathcal{Y}_i \rightarrow H_i$ be the C -quasimediant (C, C) -quasi-isometry provided by Theorem 2.1, so \mathcal{Y}_i is a CAT(0) cube complex of dimension $\leq \nu$ and the constant C depends on N .

Now we pass to (non-rescaled!) ultralimits¹. More specifically, f has an ultralimit which is a (K, K) -quasi-isometric embedding $\hat{f}: \mathbb{R}^\nu \rightarrow \hat{\mathcal{X}}$, for some ultralimit $\hat{\mathcal{X}}$ of \mathcal{X} . It is easily deduced from Corollary 2.15 that $\hat{\mathcal{X}}$ 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 -quasimediant (C, C) -quasi-isometry $\hat{c}: \hat{\mathcal{Y}} \rightarrow \hat{\mathcal{X}}$ so that the image of \hat{f} lies in the L -neighborhood of $\text{im}(\hat{c})$.

By a theorem of Huang — Theorem 1.1 of [Hua14b] — there exist n -dimensional cubical orthants O_1, \dots, O_k in $\hat{\mathcal{Y}}$ so that $d_{haus}(\hat{f}(\mathbb{R}^\nu), \hat{c}(\cup_{j=1}^k O_j)) < \infty$. Moreover, $\hat{c}(O_j)$ lies within finite Hausdorff distance of $\hat{f}(O'_j)$ for some $O'_j \subseteq \mathbb{R}^\nu$. Hence, $Q_j = f(O'_j)$ is the image of a C' -quasimediant (C', C') -quasi-isometric embedding. Thus, by Lemma 4.6, it lies within finite Hausdorff distance of a standard orthant. The Q_i are as required. \square

4.4. Controlled number of orthants. We now improve Theorem 4.13, by showing that the number of standard orthants required can be bounded in terms of the quasi-isometry constants:

Theorem 4.14 (Bounding the number of orthants). *Let \mathcal{X} be an asymphoric HHS of rank ν . For every K there exists N so that the following holds. Let $f: \mathbb{R}^\nu \rightarrow \mathcal{X}$ be a (K, K) -quasi-isometric embedding. Then there exist standard orthants $Q_i \subseteq \mathcal{X}$, $i = 1, \dots, N$, so that $d_{haus}(f(\mathbb{R}^\nu), \cup_{i=1}^N Q_i) < \infty$.*

The idea of the proof is as follows. First, by the above we have that $f(\mathbb{R}^\nu)$ lies Hausdorff-close to a finite union of standard orthants O_1, \dots, O_k . Now, each O_i makes a definite contribution to the volume growth in the quasiflat $f(\mathbb{R}^\nu)$, and this growth is in turn bounded by the quasi-isometry constants. So, k must be bounded. This is formalized in Proposition 4.16.

First, we need the following lemma, which 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 4.15. *For every $K, n \geq 1$ there exists C so that the following holds. Let $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a (K, K) -coarsely Lipschitz proper map. Then $d_{haus}(f(\mathbb{R}^n), \mathbb{R}^n) \leq C$.*

Proof. We actually show that if $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ 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 $\bar{f}: \overline{\mathbb{R}^n} \rightarrow \overline{\mathbb{R}^n}$ 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

¹If \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 $\hat{\mathcal{X}}$ is (much) bigger than \mathcal{X} .

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 \bar{f} must be surjective, for otherwise the Borsuk-Ulam theorem would force the existence of such pair of antipodal points. Since \bar{f} is surjective, then so is f , as required. \square

Proposition 4.16 (Volume growth). *For every K there exists N so that the following holds. Let $F: \mathbb{R}^\nu \rightarrow \mathcal{X}$ 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 $d_{\text{haus}}(O_i, O_j) = \infty$ when $i \neq j$, then $k \leq N$.*

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 $D = d_{\text{haus}}(F(\mathbb{R}^\nu), \bigcup_{i=1}^k O_i)$. By Lemma 4.11, 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 , $d(O'_i, O'_j) \geq 2D + 1$. 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.

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 $K' = K'(K, \mathcal{X})$, and it is a quasi-isometric embedding with constant depending on K, \mathcal{X} , and D (this dependence on D is the reason why we need Lemma 4.15 dealing with proper maps). Up to increasing K' , we can further assume that there is a (K', K') -quasi-isometry from O'_i to an orthant in \mathbb{R}^ν so that the composition of g_i and the quasi-isometry is also (K', K') -coarsely Lipschitz.

Notice that for each R and i , there exists a sub-orthant $O_i^R \subset O'_i$ so that if $x \in A_i$ has $g_i(x) \in O_i^R$, then $B_R(x) \subseteq A_i$.

Let C be as in Lemma 4.15 for K' , and set $C_1 = K'C + (K')^2$. Since the orthants O_i^R are quasi-geodesic spaces with constant depending on \mathcal{X} only, up to increasing C we can assume the following. Suppose that we have a subset $A \subseteq O_i^R$, for some R, i , with the property that $O_i^R \not\subseteq N_{C_1}(A)$. Then there exists $x \in O_i^R$ so that $d(x, A) \leq 2C_1 - 1$ but $d(x, A) > C_1$.

A further sub-orthant: We claim that for each i , there is a sub-orthant $O''_i \subset O'_i$ with the property that $O''_i \subset N_{C_1}(g_i(A_i)) \cap O'_i$.

Let $n \in \mathbb{N}$. Let O_i^n be the sub-orthant of O'_i defined above, which has the property that for all $x \in A_i$ with $g_i(x) \in O_i^n$, we have $B_n(x) \subset A_i$. If the sub-orthant O''_i with the claimed property does not exist, then, for each n , there exist $p_n \in A_i$ and $x_n \in O_i^n$ such that the following hold:

- $g_i(p_n) \in O_i^n$;
- $d(x_n, g_i(p_n)) \leq 2C_1$;
- $d(x_n, g_i(A_i)) > C_1$.

In fact, we can choose any $x_n \in O_i^n$ with $d(x_n, g_i(A_i)) \leq 2C_1 - 1$ but $d(x_n, g_i(A_i)) > C_1$, and then pick $p_n \in A_i$ “nearly witnessing” the first inequality, meaning p_n so that $d(x_n, g_i(p_n)) \leq 2C_1$.

Now, consider the (non-rescaled!) ultralimit \mathcal{R} of \mathbb{R}^ν with observation point (p_n) , which is isometric to \mathbb{R}^ν . The process of taking ultralimits induces a (K', K') -coarsely Lipschitz map f from \mathcal{R} to an ultralimit of O'_i with observation point (x_n) . Moreover, f is proper since g_i is a quasi-isometric embedding. Being an ultralimit of orthants, this ultralimit admits a quasi-isometry, h , to a subspace of \mathbb{R}^ν , with constants depending only on \mathcal{X} (actually, the ultralimit of the orthants is quasi-isometric to \mathbb{R}^ν , but we do not need this). In fact, by the choice of K' , we can choose a (K', K') -quasi-isometry h as above in such a way that $h \circ f$ is (K', K') -coarsely Lipschitz, and notice that $h \circ f$ is still proper. However, by construction the

map f is not C_1 -coarsely surjective, and thus $h \circ f$ is not C -coarsely surjective, contradicting Lemma 4.15 and thus verifying the claim.

Conclusion: 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 $\beta_R \leq tR^\nu$. Let $C' = C'(C, \nu, K)$ 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 $d(p, F(q)) \leq C'$. Distinct p yield distinct q . Moreover, $|N_i \cap B_R(F(0))| \geq t'R^\nu$ for all sufficiently large R and some $t' = t'(C', \mathcal{X})$. Since the A_i are disjoint, we have $\beta_R \geq kt'R^\nu$ for all sufficiently large R . Hence $k \leq t/t'$, and we are done. \square

Proof of Theorem 4.14. By Theorem 4.13, the image of F lies at finite Hausdorff distance from a union of orthants $\bigcup_{i=1}^k O_i$. We can assume that $d_{\text{haus}}(O_i, O_j) = \infty$ when $i \neq j$; indeed, if not, then we can drop O_i or O_j from the collection without affecting the conclusion. Hence, $k \leq N$, for N as in Proposition 4.16. \square

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 4.17. *For every K, N there exists L so that the following holds. Let $F: \mathbb{R}^\nu \rightarrow \mathcal{X}$ 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 $|A| \leq 2N$. For L as in Proposition 3.5, there exist arbitrarily large balls B' in \mathbb{R}^ν 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}^ν , as required. \square

Corollary 4.18. *For each K there exists L, N so that the following holds. Let $F: \mathbb{R}^\nu \rightarrow \mathcal{X}$ be a (K, K) -quasi-isometric embedding. Then there exist standard orthants O_1, \dots, O_N so that $F \subset \mathcal{N}_L(H_\theta(\bigcup_{i=1}^N O_i))$.*

Proof. Follows immediately from Theorem 4.14 and Lemma 4.17. \square

5. INDUCED MAPS ON HINGES: MAPPING CLASS GROUP RIGIDITY

Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. We have in mind the case where \mathcal{X} is the Cayley graph of the mapping class group of a finite-type surface, equipped with the HHS structure from [BHS19, Section 11].

In this section, we provide a new proof of quasi-isometric rigidity of mapping class groups. More generally, we study intersection patterns of quasiflats in \mathcal{X} and, under favorable conditions, extract suitable ‘‘combinatorial data’’ from it.

In the rest of this section, we will abstract from the mapping class group to the greatest extent permitted by our methods. We will need $(\mathcal{X}, \mathfrak{S})$ to be asymphoric, which, as previously noted, is a weak assumption. We will also impose three additional, more restrictive, assumptions on $(\mathcal{X}, \mathfrak{S})$, which are satisfied by the standard HHS structure on the mapping class group. First, we introduce a few relevant definitions and state the additional assumptions. Then, we discuss the generality in which these assumptions hold.

The next definition describes those subsets of \mathfrak{S} which give rise to standard flats (as defined in Definition 4.1).

Definition 5.1 (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 $\text{diam}(CU_i) = \infty$ for all $i \leq \nu$.

For each $U \in \mathfrak{S}$, we let $\partial\mathcal{C}U$ denote the Gromov boundary of $\mathcal{C}U$.

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 5.2 (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 5.3 (Ray associated to a hinge). For any $\mu \geq 0$, a μ -ray associated to a hinge $\sigma = (U, p)$ is a μ -hierarchy path \mathfrak{h}_σ so that $\pi_U(\mathfrak{h}_\sigma)$ is a quasigeodesic ray representing p and so that $\text{diam}(\pi_V(\mathfrak{h}_\sigma)) \leq \mu$ for $V \neq U$.

Remark 5.4. Any two candidates for \mathfrak{h}_σ lie at finite Hausdorff distance, so for our purposes an arbitrary choice is fine. If $\sigma \neq \sigma' \in \mathbf{Hinge}(\mathfrak{S})$, then $d_{\text{haus}}(\mathfrak{h}_\sigma, \mathfrak{h}_{\sigma'}) = \infty$.

Remark 5.5. 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 for which the product regions P_U can be taken to be subgroups. This is the case for all naturally occurring hierarchically hyperbolic structures on groups of which we are aware. However, there are some pathological structures, even on a free group, where the assumption fails.

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

Remark 5.6. In what follows, we could replace Assumption 1 with: for each $U \in \mathfrak{S}$ which is the first coordinate of some hinge, $|\partial\mathcal{C}U| \geq 2$. 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 $\{U, V\} = \mathcal{U}_1 \cap \mathcal{U}_2$.

Remark. The three preceding assumptions are, taken together, fairly restrictive. The first, as we said, is very general and holds for all “interesting” HHGs, including: mapping class groups, 3-manifolds groups which are HHG, all groups acting geometrically on CAT(0) cube complexes with *factor systems* (see [BHS17b, HS18]) (a class which includes all compact special groups in the sense of Haglund–Wise [HW08], and in particular all right-angled Artin and Coxeter groups), etc. More generally, this first assumption also holds for a number of interesting HHSs as well. These include Teichmüller spaces with the Weil-Petersson metric. On the other hand, this condition fails to hold for the HHS structure on a Teichmüller space

endowed with the Teichmüller metric, since in such structure certain \mathcal{CU} are (isometric to) horoballs in the hyperbolic plane, and thus have a single point as their boundary.

To see why the second condition is more restrictive, consider a right-angled Artin group A_Γ presented by a finite simplicial graph Γ . There are two “standard” HHS structures (see [BHS17b, Section 8] for more details), but for our purposes, we take the one described in the introduction. The second condition implies that for each vertex $v \in \Gamma$ that is contained in a maximal clique, there are two maximal cliques whose intersection is v . One can articulate a similar combinatorial condition on right-angled Coxeter groups. So, for example, the results in this section do not immediately improve upon, or even recover, Huang’s results on quasi-isometric rigidity for right-angled Artin groups.

The third condition in the right angled Artin group case can similarly be interpreted as a combinatorial constraint on the intersection pattern of cliques in the presentation graph.

In the following theorem we show that, under the additional assumptions stated above, quasi-isometries between HHSs naturally induce (orthogonality-preserving) bijections between corresponding sets of hinges. We think of such bijections as “combinatorial data” that we extract from the quasi-isometry. The proof relies on studying coarse-intersection patterns of orthants.

Theorem 5.7. *Let $(\mathcal{X}, \mathfrak{G}), (\mathcal{Y}, \mathfrak{T})$ be asymphoric HHS satisfying assumptions (1), (2) and (3). For any quasi-isometry $f: \mathcal{X} \rightarrow \mathcal{Y}$, there exists a bijection $f^\# : \mathbf{Hinge}(\mathfrak{G}) \rightarrow \mathbf{Hinge}(\mathfrak{T})$ satisfying:*

- $f^\#$ preserves orthogonality of hinges;
- for all $\sigma \in \mathbf{Hinge}(\mathfrak{G})$, we have $d_{\text{haus}}(\mathfrak{h}_{f^\#(\sigma)}, f(\mathfrak{h}_\sigma)) < \infty$.

Remark 5.8. Under suitable conditions, we expect that there exists an analogue of Theorem 5.7 in which hinges are replaced by sets of pairs $\{(U_i, p_i)\}$, where $\{U_i\}_i$ is a pairwise orthogonal set and $p_i \in \partial \mathcal{CU}_i$. 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 ≥ 2 , as in [FLS15].

Proof of Theorem 5.7. Let $\sigma = (U, p) \in \mathbf{Hinge}(\mathfrak{G})$.

How we will define $f^\#$: We will produce a hinge σ' so that $d_{\text{haus}}(\mathfrak{h}_{\sigma'}, f(\mathfrak{h}_\sigma)) < \infty$. Remark 5.4 implies that σ' is uniquely determined by this property, so we can set $f^\#(\sigma) = \sigma'$. To see that this is a bijection, let $\tilde{f}: \mathcal{Y} \rightarrow \mathcal{X}$ be a quasi-inverse of f . Then $d_{\text{haus}}(\tilde{f}(\mathfrak{h}_{\sigma'}), \mathfrak{h}_\sigma) < \infty$, so we can define an inverse for $f^\#$ in the same way.

Choosing σ' : Since (U, p) is a hinge, U is in a complete support set.

Notice that, by Assumption 1, for any complete support set $\{U_i\}_i$ we have $|\partial \mathcal{CU}_i| \geq 2$ for each i , and hence there exists a standard flat \mathcal{F} with support $\{U_i\}_i$.

In view of this, Assumption 2 provides two standard flats $\mathcal{F}_1, \mathcal{F}_2$, the intersection of whose support sets is $\{U\}$. Furthermore, we claim that we can arrange that $\mathcal{F}_1 \tilde{\cap} \mathcal{F}_2$ is coarsely a line and coarsely contains \mathfrak{h}_σ . This can be done as follows. Consider any hinge (U, q) with $q \neq p$ (which exists by Assumption 1). Assumption 2 provides complete support sets $\{V_j\}, \{V'_j\}$ whose intersection is $\{U\}$. Label so that $V_1 = V'_1 = U$. So, by choosing distinct points $a_j, b_j \in \partial \mathcal{CV}_j$ and $a'_j, b'_j \in \partial \mathcal{CV}'_j$ in such a way that $a_1 = a'_1 = p$ and $b_1 = b'_1 = q$, we obtain standard flats $\mathcal{F}_1, \mathcal{F}_2$ with the given support sets and, by Lemma 4.12, coarse intersection which is either bounded, a standard 1-orthant, or a standard 1-flat supported on U . Both flats coarsely contain a standard 1-flat (with “limit points” p and q), so the last case must hold. Moreover, the aforementioned standard 1-flat coarsely contains \mathfrak{h}_σ by Remark 5.4.

By Theorem 4.13 (Quasiflats Theorem), $f(\mathcal{F}_1)$ 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 coarsely a finite union of coarse intersections of pairs of standard orthants. Indeed, $f(\mathcal{F}_1)$ coarsely coincides with $\bigcup_s O_s$ and $f(\mathcal{F}_2)$ coarsely coincides with $\bigcup_t O'_t$, where O_s, O'_t are standard orthants provided by Theorem 4.13. Hence $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ coarsely coincides with $\bigcup_{s,t} O_s \tilde{\cap} O'_t$.
- $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is coarsely \mathbb{R} , because $\mathcal{F}_1 \tilde{\cap} \mathcal{F}_2$ was coarsely \mathbb{R} .
- $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ coarsely contains $f(\mathfrak{h}_\sigma)$, because $\mathcal{F}_1 \tilde{\cap} \mathcal{F}_2$ coarsely contained \mathfrak{h}_σ .

By Lemma 4.11 and the first of the above properties, $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is coarsely the finite union of standard k -orthants, which arise 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{CV} . 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 and Lemma 4.12 provide a standard 2-flat, \mathcal{F} , coarsely containing \mathfrak{h}_σ 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 4.13, $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is coarsely the union of finitely many coarse intersections of pairs of standard orthants. Lemma 4.11 shows that each of these intersections is coarsely a standard k -orthant for $k \geq 2$. Since $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is a 2-dimensional quasiflat, we can discard any of the above intersections which is coarsely a 0-orthant or 1-orthant. In other words, $f(\mathcal{F}_1) \tilde{\cap} f(\mathcal{F}_2)$ is coarsely the union of disjoint standard 2-orthants O_0, \dots, 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 .

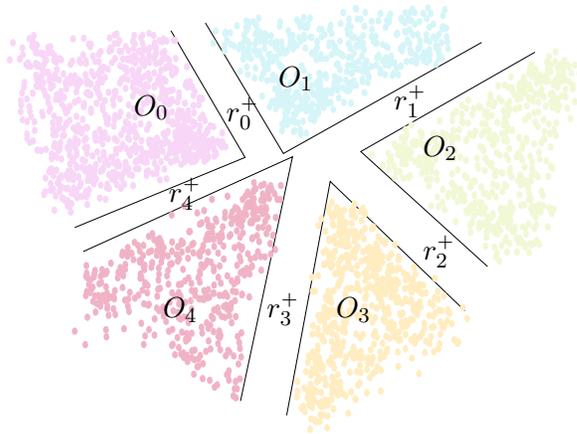


FIGURE 5. The 2-orthants O_0, \dots, O_t and the cyclic ordering of their coordinate rays (up to coarse coincidence).

Now, as shown in Figure 5, we can cyclically order the coordinate rays in O_0, \dots, 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: $r_0^+, r_1^+, \dots, r_{t-1}^+$.

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, and bijectivity of f^\sharp .) 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 4.11. \square

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. The content of the lemma is that if a quasi-isometry matches the “combinatorial data” of a standard flat to the data of another standard flat, then it maps the former flat within uniform distance of the latter.

Lemma 5.9 (Flats go to flats). *Let $(\mathcal{X}, \mathfrak{S}), (\mathcal{Y}, \mathfrak{T})$ be asymptotic 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 $k = 1, \dots, n$ we have $f^\sharp(U_k, p_k^\pm) = (V_k, q_k^\pm)$. Then, $d_{\text{haus}}(f(\mathcal{F}_{\{(U_i, p_i^\pm)\}}), \mathcal{F}_{\{(V_j, q_j^\pm)\}}) \leq C$.*

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 4.17. The other containment follows by applying the same argument to a quasi-inverse of f . \square

5.2. Mapping class groups. We now use Theorem 5.7 to provide a new proof of quasi-isometric rigidity of mapping class groups. Like all proofs of quasi-isometric rigidity for mapping class groups, the goal of our proof is to prove that any quasi-isometry of the mapping class group induces a simplicial automorphism of \mathcal{CS} , at which point we can apply Ivanov’s theorem [Iva97] to conclude that the automorphism is induced by an element of the mapping class group. Using our quasiflats theorem we can readily convert the geometric information of a quasi-isometry to combinatorial information about the structure of standard flats. Then, via Theorem 5.7, from the combinatorial structure of quasiflats we can extract an induced map on certain coordinate directions in the standard flats. In the mapping class group setting, these directions correspond to Dehn twist directions, thus giving us the automorphism of the curve graph which is needed to apply Ivanov’s theorem.

Theorem 5.10. [BKMM12] *Let \mathcal{X} be the mapping class group of a non-sporadic surface S . Then for any K there exists L so that: for each (K, K) -quasi-isometry $f: \mathcal{X} \rightarrow \mathcal{X}$ 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{CY} are curve complexes. (For details on the structure, see [BHS19, 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 5.7 are clearly satisfied.

Consider any quasi-isometry $f: \mathcal{X} \rightarrow \mathcal{X}$. 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 6, where, if σ is (U, p^+) , then σ' is (U, p^-) , where $\partial\mathcal{C}U = \{p^\pm\}$.

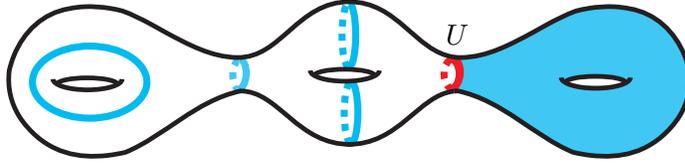


FIGURE 6. This figure shows a complete support set, consisting of five annuli and one once-punctured torus. This is the only complete support set containing all the subsurfaces except the annulus about the boundary of the once-punctured torus, denoted U in the figure. In this sense, U is non-replaceable.

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 for some maximal set \mathcal{H} of pairwise orthogonal hinges containing (U, p^+) , the hinge (U, p^-) is the only hinge such that $\mathcal{H} - \{(U, p^+)\} \cup \{(U, p^-)\}$ is a maximal set of pairwise orthogonal hinges. In this sense, the annulus U is “non-replaceable”. We write $V = f^*(U)$. Notice that Lemma 5.9 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{CS} , 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{CS} 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. \square

6. FACTORED SPACES

In this section we show that, under certain circumstances, quasi-isometries between HHSs descend to quasi-isometries between some of their “factored” spaces, which are spaces obtained by coning off a collection of standard product regions. These factored spaces are HHSs themselves and their complexity is lower than the complexity of the original HHS. Hence, studying induced quasi-isometries on factored spaces can be part of an inductive procedure for studying quasi-isometries of the original space (see also the introduction).

Notation 6.1. Given $\mathfrak{U} \subseteq \mathfrak{S}$, let \mathfrak{U}^\square 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 $\hat{\mathcal{X}}$ be the factored space associated to \mathfrak{U}^\square , which is the space obtained from \mathcal{X} by coning off all F_U for $U \in \mathfrak{U}^\square$ (as described in [BHS17a, Definition 2.1]). There exists a Lipschitz factor map $q = q_\mathcal{X}: \mathcal{X} \rightarrow \hat{\mathcal{X}}$ by [BHS17a, Proposition 2.2].

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

Theorem 6.2 (Quasiflats collapse in factored spaces). *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 $f: \mathbb{R}^\nu \rightarrow \mathcal{X}$, we have $\text{diam}(q \circ f(\mathbb{R}^\nu)) \leq \Delta$.*

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

Hence, by Corollary 4.18, it suffices to prove that $\text{diam}_{\hat{\mathcal{X}}} q(\bigcup_{i=1}^N O_i) \leq C$, 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 $\text{diam}_{\hat{\mathcal{X}}}(q(O_i)) \leq C'$ 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. \square

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 (Quasi-isometries descend to factored spaces). *Let \mathcal{X}, \mathcal{Y} be asymphoric HHSs. 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 $f: \mathcal{X} \rightarrow \mathcal{Y}$ there exists a quasi-isometry $\hat{f}: \hat{\mathcal{X}} \rightarrow \hat{\mathcal{Y}}$ so that the diagram*

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

commutes.

Proof. Since $\hat{\mathcal{X}}, \hat{\mathcal{Y}}$ are just re-metrized copies of \mathcal{X}, \mathcal{Y} (see [BHS17a]), 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 $\hat{\mathcal{X}}, \hat{\mathcal{Y}}$ ([BHS17a, Definition 2.1]), we just have to verify that if x, y lie in some F_U for $U \in \mathfrak{U}_{\mathcal{X}}^{\square}$, then their images are uniformly close in $\hat{\mathcal{Y}}$. By assumption, x, y lie close to a quasiflat with uniform constant, so that the conclusion follows from Theorem 6.2. \square

We can also now prove Theorem G from the introduction.

Proof of Theorem G. In the case where \mathcal{X} is a mapping class group, we have seen that the standard HHS structure $(\mathcal{X}, \mathfrak{S})$ is asymphoric of finite rank ν equal to the complexity. Let \mathfrak{U} be as in Notation 6.1 and let $q: \mathcal{X} \rightarrow \hat{\mathcal{X}}$ be the factor map described above. Then any ν -dimensional quasiflat in \mathcal{X} has uniformly bounded image in $\hat{\mathcal{X}}$ by Theorem 6.2. Now, letting $S \in \mathfrak{S}$ be the unique \square -maximal element, we have that $\pi_U: \hat{\mathcal{X}} \rightarrow \mathcal{CS}$ is a Lipschitz map, and the theorem follows. \square

Now we turn to the lemmas.

The following lemma identifies possible distance formula terms for pairs of points each in a given orthant. Roughly, they can be of two types, each corresponding to one of the factors of the bridge as in Lemma 1.20 between the orthants.

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 $d_U(x, y) \geq \tau$ 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 4.11 (and also with a standard orthant whose support is contained in $\mathcal{U}_1 \cap \mathcal{U}_2$, thereby describing \mathcal{U}).

Let x, y, U be as in the statement. If $x, y \in O$, then by the definition of a standard orthant, either $\mathfrak{d}_U(x, y)$ is uniformly bounded or $U \sqsubseteq U'$ for some $U' \in \mathcal{U}_1$. If $U' \in \mathcal{U}_1 \cap \mathcal{U}_2$, we are done; otherwise, $U' \perp V$ for all $V \in \mathcal{U}$ by the definition of a standard orthant. An identical argument works if $x, y \in O'$.

So, assume that $x \in O', y \in O$. If U is nested in some element of \mathcal{U}_i , for some $i \in \{1, 2\}$, then either U is nested into some element of $\mathcal{U}_1 \cap \mathcal{U}_2$, or U is orthogonal to every element of \mathcal{U} by the definition of a standard orthant. Hence, suppose that U is not nested into any element of \mathcal{U}_1 or \mathcal{U}_2 . In particular, $\pi_U(O)$ and $\pi_U(O')$ have uniformly bounded diameter, so $\mathfrak{d}_U(x, y) \asymp \mathfrak{d}_U(O, O')$. Therefore, provided τ is sufficiently large, $\mathfrak{d}_U(x, y) > \tau$ implies that $\mathfrak{d}_U(O, O')$ is large compared to the constant K_2 from Lemma 1.20, so that conclusion (5) of the same lemma (applied with any $p \in \mathfrak{g}_O(O')$, $t_1 = a$, and $t_2 = b$) shows that U is orthogonal to each element of \mathcal{U} . \square

If the coarse intersection $O \tilde{\cap} O'$ is a codimension-1 sub-orthant, then $q(O \cup O')$ is *uniformly* bounded:

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 $\text{diam}_{\hat{\chi}}(q(O \cup O')) \leq C$.*

Proof. Let $x \in O, y \in O'$. Let $\mathcal{M} = \{U \in \mathfrak{S} : \mathfrak{d}_U(x, y) \geq \tau\}$. 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 $\mathfrak{d}_U(q(x), q(y)) \leq \tau$ for all $U \in \mathfrak{S} - \mathcal{M}$, so $q(x)$ is uniformly close to $q(y)$ by the uniqueness axiom. \square

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 $\mathfrak{d}_{\text{haus}}(O_i, O_j) = \infty$ for $i \neq j$. Then for each pair of distinct orthants O_j, O_k there exists a sequence $j = j_0, \dots, j_l = k$ 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 \mathcal{F} of \mathbb{R}^ν filled by bilipschitz copies \mathbf{O}_i of $[0, \infty)^\nu$. The intersections of the \mathbf{O}_i have some basic properties:

Lemma 6.7.

- (1) *The intersection of \mathbf{O}_i and \mathbf{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 \mathbf{O}_i and \mathbf{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 4.11). We now show the following, which implies both statements: if the ultralimits \mathbf{A}, \mathbf{B} of uniformly hierarchically quasiconvex sets have non-empty intersection, then their intersection is the ultralimit $\mathfrak{g}_{\mathbf{A}}(\mathbf{B})$ of the gates. By Lemma 1.20.(3), $\mathfrak{g}_{\mathbf{A}}(\mathbf{B})$ is contained in $\mathbf{A} \cap \mathbf{B}$ (this uses $\mathfrak{d}(\mathbf{A}, \mathbf{B}) = 0$). Lemma 1.20.(6) implies that the other containment holds. \square

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

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

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

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

We first claim that for any $O \in \mathcal{Y}$ and any closed $O' \subset O$, $\mathcal{F} - O'$ is path-connected and $H_1(\mathcal{F} - O') = 0$. We use the fact that, for A, B closed homeomorphic subsets of \mathbb{R}^ν , 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 $H_1(\mathcal{F} - O') = 0$ follows from the fact that $H_1(\mathcal{F} - O) = 0$, since a 1-cycle in $\mathcal{F} - O'$ is homologous to one in $\mathcal{F} - O$ by, for example, a transversality argument. The same holds for $H_0(\mathcal{F} - O')$.

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:

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

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

We now finish the proof of Proposition 6.6.

Let $\mathbf{O}_j, \mathbf{O}_k$ be orthants. We will now produce a sequence $\mathbf{O}_j = \mathbf{O}_{j_0}, \dots, \mathbf{O}_{j_l} = \mathbf{O}_k$ of orthants so that $t(j_i, j_{i+1}) = \nu - 1$ for $0 \leq i \leq l - 1$. Choose $\mathbf{x} \in \text{Int}(\mathbf{O}_i), \mathbf{y} \in \text{Int}(\mathbf{O}_j)$ and let $\sigma: [0, 1] \rightarrow \mathcal{F} - Y$ be a path joining them, which is provided by Lemma 6.8. Let t_0 be the maximal t so that $\sigma(t) \in \mathbf{O}_j$. If $t_0 = 1$, then we take $l = 0$. Otherwise, there exists $\mathbf{O}_{j_1} \neq \mathbf{O}_j$ so that $\mathbf{O}_j \cap \mathbf{O}_{j_1}$ 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. \square

REFERENCES

- [BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math. Inst. Hautes Études Sci.*, 122:1–64, 2015.
- [BBF19] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Proper actions on finite products of quasi-trees. *arXiv 1905.10813*, 2019.
- [BC11] Jason Behrstock and Ruth Charney. Divergence and quasimorphisms of right-angled Artin groups. *Mathematische Annalen*, pages 1–18, 2011.
- [BDM09] J. Behrstock, C. Druţu, and L. Mosher. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. *Math. Ann.*, 344:543–595, 2009.
- [Beh06] J. Behrstock. Asymptotic geometry of the mapping class group and Teichmüller space. *Geometry & Topology*, 10:1523–1578, 2006.
- [BF06] J. Brock and B. Farb. Curvature and rank of Teichmüller space. *American Jour. Math.*, 128(1):1–22, 2006.
- [BHS17a] J. Behrstock, M.F. Hagen, and A. Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. London Math. Soc.*, 114(5):890–926, 2017.
- [BHS17b] J. Behrstock, M.F. Hagen, and A. Sisto. Hierarchically hyperbolic spaces I: curve complexes for cubical groups. *Geom. Topol.*, 21:1731–1804, 2017.
- [BHS19] J. Behrstock, M.F. Hagen, and A. Sisto. Hierarchically hyperbolic spaces II: combination theorems and the distance formula. *Pacific J. Math.*, 299:257–338, 2019.
- [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.
- [BM08a] J. Behrstock and Y. Minsky. Dimension and rank for mapping class groups. *Ann. of Math.*, 167:1055–1077, 2008.
- [BM08b] 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.
- [Bow18a] Brian H Bowditch. Convex hulls in coarse median spaces. *Preprint*, 2018.
- [Bow18b] Brian H. Bowditch. Large-scale rigidity properties of the mapping class groups. *Pacific J. Math.*, 293(1):1–73, 2018.
- [BR17] Federico Berlai and Bruno Robbio. A refined combination theorem for hierarchically hyperbolic spaces. *In preparation*, 2017.
- [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.
- [Bro03] J. Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *Jour. AMS*, 16(3):495–535 (electronic), 2003.
- [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.
- [DK18] Cornelia Druțu and Michael Kapovich. *Geometric group theory*, volume 63 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2018. With an appendix by Bogdan Nica.
- [Dol93] Albrecht Dold. A simple proof of the Jordan-Alexander complement theorem. *Amer. Math. Monthly*, 100(9):856–857, 1993.
- [Dru02] Cornelia Druțu. Quasi-isometry invariants and asymptotic cones. *International Journal of Algebra and Computation*, 12(01n02):99–135, 2002.
- [DS05] C. Druțu and M. Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005. With an appendix by Denis Osin and Mark Sapir.
- [EF97] A. Eskin and B. Farb. Quasi-flats and rigidity in higher rank symmetric spaces. *Jour. AMS*, 10(3):653–692, 1997.
- [EMR17] Alex Eskin, Howard Masur, and Kasra Rafi. Large-scale rank of Teichmüller space. *Duke Math. J.*, 166(8):1517–1572, 2017.
- [FLS15] Roberto Frigerio, Jean-François Lafont, and Alessandro Sisto. Rigidity of high dimensional graph manifolds. *Astérisque*, (372):xxi+177, 2015.
- [Gro87a] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [Gro87b] 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.
- [HS18] Mark F Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. *Isr. Jour. Math.*, To appear., 2018. 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.
- [HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.
- [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.

- [KK14] Sang-Hyun Kim and Thomas Koberda. The geometry of the curve graph of a right-angled Artin group. *International Journal of Algebra and Computation*, 24(02):121–169, 2014.
- [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.
- [MM00] H. Masur and Y. Minsky. Geometry of the complex of curves II: Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.
- [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.
- [RST18] Jacob Russell, Davide Spriano, and Hung Cong Tran. Convexity in hierarchically hyperbolic spaces. *arXiv preprint arXiv:1809.09303*, 2018.
- [Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
- [Sch95] R. Schwartz. The quasi-isometry classification of rank one lattices. *Inst. Hautes Études Sci. Publ. Math.*, 82:133–168, 1995.
- [Tao] Terence Tao. Nonstandard analysis as a completion of standard analysis. <https://terrytao.wordpress.com/2010/11/27/nonstandard-analysis-as-a-completion-of-standard-analysis>.
- [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
Current address: School of Mathematics, University of Bristol, Bristol, UK
E-mail address: `markfhagen@gmail.com`

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