

HIERARCHICALLY HYPERBOLIC SPACES II: COMBINATION THEOREMS AND THE DISTANCE FORMULA

JASON BEHRSTOCK, MARK F. HAGEN, AND ALESSANDRO SISTO

ABSTRACT. We introduce a number of tools for finding and studying *hierarchically hyperbolic spaces* (HHS), a rich class of spaces including mapping class groups of surfaces, Teichmüller space with either the Teichmüller or Weil-Petersson metrics, right-angled Artin groups, and the universal cover of any compact special cube complex. We begin by introducing a streamlined set of axioms defining an HHS. We prove that all HHS satisfy a Masur-Minsky-style distance formula, thereby obtaining a new proof of the distance formula in the mapping class group without relying on the Masur-Minsky hierarchy machinery. We then study examples of HHS; for instance, we prove that when M is a closed irreducible 3-manifold then $\pi_1 M$ is an HHS if and only if it is neither *Nil* nor *Sol*. We establish this by proving a general combination theorem for trees of HHS (and graphs of HH groups). We also introduce a notion of “hierarchical quasiconvexity”, which in the study of HHS is analogous to the role played by quasiconvexity in the study of Gromov-hyperbolic spaces.

CONTENTS

Introduction	1
1. The main definition and background on hierarchically hyperbolic spaces	9
2. Tools for studying hierarchically hyperbolic spaces	16
3. Realization of consistent tuples	19
4. Hierarchy paths and the distance formula	22
5. Hierarchical quasiconvexity and gates	29
6. Hulls	33
7. The coarse median property	37
8. Combination theorems for hierarchically hyperbolic spaces	39
9. Hyperbolicity relative to HHGs	52
10. Hierarchical hyperbolicity of 3-manifold groups	54
11. A new proof of the distance formula for mapping class groups	56
References	58

INTRODUCTION

One of the most remarkable aspects of the theory of mapping class groups of surfaces is that the coarse geometry of the mapping class group, $\mathcal{MCG}(S)$, can be fully reconstructed from its shadows on a collection of hyperbolic spaces — namely the curve graphs of subsurfaces of the underlying surface. Each subsurface of the surface S is equipped with a hyperbolic curve graph and a projection, the *subsurface projection*, to this graph from $\mathcal{MCG}(S)$;

Behrstock was supported as a Simons Fellow.

Hagen was supported by NSF Grant Number 1045119 and by EPSRC.

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

there are also projections between the various curve graphs. The powerful Masur–Minsky distance formula [MM00] shows that the distance between points of $\mathcal{MCG}(S)$ is coarsely the sum over all subsurfaces of the distances between the projections of these points to the various curve graphs. Meanwhile, the consistency/realization theorem [BKMM12] tells us that tuples with coordinates in the different curve graphs that obey “consistency” conditions characteristic of images of actual points in $\mathcal{MCG}(S)$ are, coarsely, images of points in $\mathcal{MCG}(S)$. Finally, any two points in $\mathcal{MCG}(S)$ are joined by a uniform-quality quasigeodesic projecting to a uniform unparameterized quasigeodesic in each curve graph — a *hierarchy path* [MM00].

It is perhaps surprising that analogous behavior should appear in CAT(0) cube complexes, since the mapping class group cannot act properly on such complexes, c.f., [Bri10, Hag07, KL96]. However, mapping class groups enjoy several properties reminiscent of non-positively/negatively curved spaces, including: automaticity (and, thus, quadratic Dehn function) [Mos95], having many quasimorphisms [BF02], super-linear divergence [Beh06], etc. Mapping class groups also exhibit coarse versions of some features of CAT(0) cube complexes, including coarse centroids/medians [BM11] and, more generally, a local coarse structure of a cube complex as made precise in [Bow13], applications to embeddings in trees, [BDS11], etc. Accordingly, it is natural to seek a common thread joining these important classes of groups and spaces.

In [Hag14] it was shown that, for an arbitrary CAT(0) cube complex \mathcal{X} , the intersection-graph of the hyperplane carriers — the *contact graph* — is hyperbolic, and in fact quasi-isometric to a tree. This object seems at first glance quite different from the curve graph (which records, after all, *non*-intersection), but there are a number of reasons this is quite natural, two of which we now mention. First, the curve graph can be realized as a coarse intersection graph of product regions in \mathcal{MCG} . Second, the contact graph is closely related to the intersection graph of the hyperplanes themselves; when \mathcal{X} is the universal cover of the Salvetti complex of a right-angled Artin group, the latter graph records commutation of conjugates of generators, just as the curve graph records commutation of Dehn twists.

The cube complex \mathcal{X} coarsely projects to its contact graph. Moreover, using disc diagram techniques, it is not hard to show that any two 0-cubes in a CAT(0) cube complex are joined by a combinatorial geodesic projecting to a geodesic in the contact graph [BHS17a]. This observation — that CAT(0) cube complexes have “hierarchy paths” with very strong properties — motivated a search for an analogue of the theory of curve graphs and subsurface projections in the world of CAT(0) cube complexes. This was largely achieved in [BHS17a], where a theory completely analogous to the mapping class group theory was constructed for a wide class of CAT(0) cube complexes, with (a variant of) the contact graph playing the role of the curve graph. (Results of this type for right-angled Artin groups, using the *extension graph*, were obtained by Kim-Koberda in [KK14]; see [BHS17a] for a comparison of the two approaches.)

These results motivated us to define a notion of “spaces with distance formulae”, which we did in [BHS17a], by introducing the class of *hierarchically hyperbolic spaces (HHS)* to provide a framework for studying many groups and spaces which arise naturally in geometric group theory, including mapping class groups and virtually special groups, and to provide a notion of “coarse nonpositive curvature” which is quasi-isometry invariant while still yielding some of those properties available via local geometry in the classical setting of nonpositively-curved spaces.

As mentioned above, the three most salient features of hierarchically hyperbolic spaces are: the distance formula, the realization theorem, and the existence of hierarchy paths. In the treatment given in [BHS17a], these attributes are part of the definition of a hierarchically

hyperbolic space. This is somewhat unsatisfactory since, in the mapping class group and cubical settings, proving these theorems requires serious work.

In this paper, we show that although the definition of hierarchically hyperbolic space previously introduced identifies the right class of spaces, there exist a streamlined set of axioms for that class of spaces which are much easier to verify in practice than those presented in [BHS17a, Section 13] and which don't require assuming a distance formula, realization theorem, or the existence of hierarchy paths. Thus, a significant portion of this paper is devoted to proving that those results can be derived from the simplified axioms we introduce here. Along the way, we obtain a new, simplified proof of the actual Masur-Minsky distance formula for the mapping class group. We then examine various geometric properties of hierarchically hyperbolic spaces and groups, including many reminiscent of the world of CAT(0) spaces and groups; for example, we show, using an argument due to Bowditch, that hierarchically hyperbolic groups have quadratic Dehn function. Finally, taking advantage of the simpler set of axioms, we prove combination theorems enabling the construction of new hierarchically hyperbolic spaces/groups from old.

The definition of a hierarchically hyperbolic space still has several parts, the details of which we postpone to Section 1. However, the idea is straightforward: a hierarchically hyperbolic space is a pair $(\mathcal{X}, \mathfrak{S})$, where \mathcal{X} is a metric space and \mathfrak{S} indexes a set of δ -hyperbolic spaces with several features (to each $U \in \mathfrak{S}$ the associated space is denoted $\mathcal{C}U$). Most notably, \mathfrak{S} is endowed with 3 mutually exclusive relations, *nesting*, *orthogonality*, and *transversality*, respectively generalizing nesting, disjointness, and overlapping of subsurfaces. For each $U \in \mathfrak{S}$, we have a coarsely Lipschitz projection $\pi_U: \mathcal{X} \rightarrow \mathcal{C}U$, and there are relative projections $\mathcal{C}U \rightarrow \mathcal{C}V$ when $U, V \in \mathfrak{S}$ are non-orthogonal. These projections are required to obey “consistency” conditions modeled on the inequality identified by Behrstock in [Beh06], as well as a version of the bounded geodesic image theorem and large link lemma of [MM00], among other conditions. A finitely generated group G is *hierarchically hyperbolic* if it can be realized as a group of HHS automorphisms (“hieromorphisms”, as defined in Section 1) so that the induced action on \mathcal{X} by uniform quasi-isometries is geometric and the action on \mathfrak{S} is cofinite. Hierarchically hyperbolic groups, endowed with word-metrics, are hierarchically hyperbolic spaces, but the converse does not appear to be true.

Combination theorems. One of the main contributions in this paper is to provide many new examples of hierarchically hyperbolic groups, thus showing that mapping class groups and various cubical complexes/groups are just two of many interesting families in this class of groups and spaces. We provide a number of combination theorems, which we will describe below. One consequence of these results is the following classification of exactly which 3-manifold groups are hierarchically hyperbolic:

Theorem 10.1 (3-manifolds are hierarchically hyperbolic). *Let M be a closed 3-manifold. Then $\pi_1(M)$ is a hierarchically hyperbolic space if and only if M does not have a Sol or Nil component in its prime decomposition.*

This result has a number of applications to the many fundamental groups of 3-manifolds which are HHS. For instance, in such cases, it follows from results in [BHS17a] that: except for \mathbb{Z}^3 , the top dimension of a quasi-flat in such a group is 2, and any such quasi-flat is locally close to a “standard flat” (this generalizes one of the main results of [KL97, Theorem 4.10]); up to finite index, \mathbb{Z} and \mathbb{Z}^2 are the only finitely generated nilpotent groups which admit quasi-isometric embeddings into $\pi_1(M)$; and, except in the degenerate case where $\pi_1(M)$ is virtually abelian, such groups are all acylindrically hyperbolic (as also shown in [MO14]).

Remark (Hierarchically hyperbolic spaces vs. hierarchically hyperbolic groups). There is an important distinction to be made between a *hierarchically hyperbolic space*, which is a metric

space \mathcal{X} equipped with a collection \mathfrak{S} of hyperbolic spaces with certain properties, and a *hierarchically hyperbolic group*, which is a group acting geometrically on a hierarchically hyperbolic space in such a way that the induced action on \mathfrak{S} is cofinite. The latter property is considerably stronger. For example, Theorem 10.1 shows that $\pi_1 M$, with any word-metric, is a hierarchically hyperbolic space, but, as we discuss in Remark 10.2, $\pi_1 M$ probably fails to be a hierarchically hyperbolic group in general; for instance we conjecture this is the case for those graph manifolds which can not be cocompactly cubulated.

In the course of proving Theorem 10.1, we establish several general combination theorems, including one about relative hyperbolicity and one about graphs of groups. The first is:

Theorem 9.1 (Hyperbolicity relative to HHGs). *Let the group G be hyperbolic relative to a finite collection \mathcal{P} of peripheral subgroups. If each $P \in \mathcal{P}$ is a hierarchically hyperbolic space, then G is a hierarchically hyperbolic space. Further, if each $P \in \mathcal{P}$ is a hierarchically hyperbolic group, then so is G .*

Another of our main results is a combination theorem, Theorem 8.6, establishing when a tree of hierarchically hyperbolic spaces is again a hierarchically hyperbolic space. In the statement below, *hierarchical quasiconvexity* is a natural generalization of both quasiconvexity in the hyperbolic setting and cubical convexity in the cubical setting, which we shall discuss in some detail shortly. The remaining conditions are technical and explained in Section 8, but are easily verified in practice.

Theorem 8.6 (Combination theorem for HHS). *Let \mathcal{T} be a tree of hierarchically hyperbolic spaces. Suppose that:*

- *edge-spaces are uniformly hierarchically quasiconvex in incident vertex spaces;*
- *each edge-map is full;*
- *\mathcal{T} has bounded supports;*
- *if e is an edge of \mathcal{T} and S_e is the \sqsubseteq -maximal element of \mathfrak{S}_e , then for all $V \in \mathfrak{S}_{e^\pm}$, the elements V and $\phi_{e^\pm}^\diamond(S_e)$ are not orthogonal in \mathfrak{S}_{e^\pm} .*

Then $\mathcal{X}(\mathcal{T})$ is hierarchically hyperbolic.

As a consequence, we obtain a set of sufficient conditions guaranteeing that a graph of hierarchically hyperbolic groups is a hierarchically hyperbolic group.

Corollary 8.24 (Combination theorem for HHG). *Let $\mathcal{G} = (\Gamma, \{G_v\}, \{G_e\}, \{\phi_e^\pm\})$ be a finite graph of hierarchically hyperbolic groups. Suppose that \mathcal{G} equivariantly satisfies the hypotheses of Theorem 8.6. Then the total group G of \mathcal{G} is a hierarchically hyperbolic group.*

Finally, we prove that products of hierarchically hyperbolic spaces admit natural hierarchically hyperbolic structures.

As mentioned earlier, we will apply the combination theorems to fundamental groups of 3-manifolds, but their applicability is broader. For example, they can be applied to fundamental groups of higher dimensional manifolds such as the ones considered in [FLS15].

The distance formula and realization. As defined in [BHS17a], the basic definition of a *hierarchically hyperbolic space* is modeled on the essential properties underlying the “hierarchy machinery” of mapping class groups. In this paper, we revisit the basic definition and provide a new, refined set of axioms; the main changes are the removal of the “distance formula” and “hierarchy path” axioms and the replacement of the “realization” axiom by a far simpler “partial realization”. These new axioms are both more fundamental and more readily verified.

An important result in mapping class groups which provides a starting point for much recent research in the field is the celebrated “distance formula” of Masur–Minsky [MM00]

which provides a way to estimate distances in the mapping class group, up to uniformly bounded additive and multiplicative distortion, via distances in the curve graphs of subsurfaces. We give a new, elementary, proof of the distance formula in the mapping class group. The first step in doing so is verifying that mapping class groups satisfy the new axioms of a hierarchically hyperbolic space. We provide elementary, simple proofs of the axioms for which elementary proofs do not exist in the literature (most notably, the uniqueness axiom); this is done in Section 11. This then combines with our proof of the following result which states that any hierarchically hyperbolic space satisfies a “distance formula” (which in the case of the mapping class group provides a new proof of the Masur–Minsky distance formula):

Theorem 4.5 (Distance formula for HHS). *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 $x, y \in \mathcal{X}$,*

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

where $d_W(x, y)$ denotes the distance in the hyperbolic space CW between the projections of x, y and $\{\{A\}\}_B = A$ if $A \geq B$ and 0 otherwise.

Moreover, we show in Theorem 4.4 that any two points in \mathcal{X} are joined by a uniform quasigeodesic γ projecting to a uniform unparameterized quasigeodesic in CU for each $U \in \mathfrak{S}$. The existence of such *hierarchy paths* was hypothesized as part of the definition of a hierarchically hyperbolic space in [BHS17a], but now it is proven as a consequence of the other axioms.

The Realization Theorem for the mapping class group was established by Behrstock–Kleiner–Minsky–Mosher in [BKMM12, Theorem 4.3]. This theorem states that given a surface S and, for each subsurface $W \subseteq S$, a point in the curve complex of W : this sequence of points arises as the projection of a point in the mapping class group (up to bounded error), whenever the curve complex elements satisfy certain pairwise “consistency conditions.” Thus the Realization Theorem provides another sense in which all of the quasi-isometry invariant geometry of the mapping class group is encoded by the projections onto the curve graphs of subsurfaces.¹ In Section 3 we show that an arbitrary hierarchically hyperbolic space satisfies a realization theorem. Given our elementary proof of the new axioms for mapping class groups in Section 11, we thus obtain a new proof of [BKMM12, Theorem 4.3].

Hulls and the coarse median property. Bowditch introduced a notion of *coarse median space* to generalize some results about median spaces to a more general setting, and, in particular, to the mapping class group [Bow13]. Bowditch observed in [Bow18] that any hierarchically hyperbolic space is a coarse median space; for completeness we provide a short proof of this result in Theorem 7.3. Using Bowditch’s results about coarse median spaces, we obtain a number of applications as corollaries. For instance, Corollary 7.9 is obtained from [Bow14a, Theorem 9.1] and says that any hierarchically hyperbolic space satisfies the Rapid Decay Property and Corollary 7.5 is obtained from [Bow13, Corollary 8.3] to show that all hierarchically hyperbolic groups are finitely presented and have quadratic Dehn functions. This provides examples of groups that are not hierarchically hyperbolic, for example:

Corollary 7.6 ($\text{Out}(F_n)$ is not an HHG). *For $n \geq 3$, the group $\text{Out}(F_n)$ is not a hierarchically hyperbolic group.*

¹In [BKMM12], the name Consistency Theorem is used to refer to the necessary and sufficient conditions for realization; since we find it useful to break up these two aspects, we refer to this half as the Realization Theorem, since anything that satisfies the consistency conditions is realized.

Indeed, $\text{Out}(F_n)$ was shown in [BV95, HM13b, BV] to have exponential Dehn function. This result is interesting as a counter-point to the well-known and fairly robust analogy between $\text{Out}(F_n)$ and the mapping class group of a surface; especially in light of the fact that $\text{Out}(F_n)$ is known to have a number of properties reminiscent of the axioms for an HHS, c.f., [BF14b, BF14a, HM13a, SS].

The coarse median property, via work of Bowditch, also implies that asymptotic cones of hierarchically hyperbolic spaces are contractible. Moreover, in Corollary 6.7, we bound the homological dimension of any asymptotic cone of a hierarchically hyperbolic space. This latter result relies on the use of *hulls* of finite sets of points in the HHS \mathcal{X} . This construction generalizes the Σ -*hull* of a finite set, constructed in the mapping class group context in [BKMM12]. (It also generalizes a special case of the ordinary combinatorial convex hull in a CAT(0) cube complex.) A key feature of these hulls is that they are coarse retracts of \mathcal{X} (see Proposition 6.3), and this plays an important role in the proof of the distance formula.

Hierarchical spaces. We also introduce the more general notion of a *hierarchical space* (HS). This is the same as a hierarchically hyperbolic space, except that we do not require the various associated spaces \mathcal{CU} , onto which we are projecting, to be hyperbolic. Although we mostly focus on HHS in this paper, a few things are worth noting. First, the realization theorem (Theorem 3.1) actually makes no use of hyperbolicity of the \mathcal{CU} , and therefore holds in the more general context of HS; see Section 3.

Second, an important subclass of the class of HS is the class of *relatively hierarchically hyperbolic spaces*, which we introduce in Section 6.2. These are hierarchical spaces where the spaces \mathcal{CU} are uniformly hyperbolic except when U is minimal with respect to the nesting relation. As their name suggests, this class includes all metrically relatively hyperbolic spaces; see Theorem 9.3. With an eye to future applications, in Section 6.2 we prove a distance formula analogous to Theorem 4.5 for relatively hierarchically hyperbolic spaces, and also establish the existence of hierarchy paths. The strategy is to build, for each pair of points x, y , in the relatively hierarchically hyperbolic space, a “hull” of x, y , which we show is hierarchically hyperbolic with uniform constants. We then apply Theorems 4.5 and 4.4.

Standard product regions and hierarchical quasiconvexity. In Section 5.1, we introduce the notion of a *hierarchically quasiconvex* subspace of a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$. In the case where \mathcal{X} is hyperbolic, this notion coincides with the usual notion of quasiconvexity. The main technically useful features of hierarchically quasiconvex subspaces generalize key features of quasiconvexity: they inherit the property of being hierarchically hyperbolic (Proposition 5.6) and one can coarsely project onto them (Lemma 5.5).

Along with the hulls discussed above, the most important examples of hierarchically quasiconvex subspaces are *standard product regions*: for each $U \in \mathfrak{S}$, one can consider the set P_U of points $x \in \mathcal{X}$ whose projection to each \mathcal{CV} is allowed to vary only if V is nested into or orthogonal to U ; otherwise, x projects to the same place in \mathcal{CV} as \mathcal{CU} does under the relative projection. The space P_U coarsely decomposes as a product, with factors corresponding to the nested and orthogonal parts. Product regions play an important role in the study of boundaries and automorphisms of hierarchically hyperbolic spaces in the forthcoming paper [DHS17], as well as in the study of quasi-boxes and quasiflats in hierarchically hyperbolic spaces carried out in [BHS17a].

Some questions and future directions. Before embarking on the discussion outlined above, we raise a few questions about hierarchically hyperbolic spaces and groups that we believe are of significant interest.

The first set of questions concern the scope of the theory, i.e., which groups and spaces are hierarchically hyperbolic and which operations preserve the class of HHS:

Question A (Cubical groups). *Let G act properly and cocompactly on a $CAT(0)$ cube complex. Is G a hierarchically hyperbolic group? Conversely, suppose that (G, \mathfrak{S}) is a hierarchically hyperbolic group; are there conditions on the elements of \mathfrak{S} which imply that G acts properly and cocompactly on a $CAT(0)$ cube complex?*²

Substantial evidence for this conjecture was provided in [BHS17a] where we established that a $CAT(0)$ cube complex \mathcal{X} containing a *factor system* is a hierarchically hyperbolic space, and the associated hyperbolic spaces are all uniform quasi-trees. (Roughly speaking, \mathcal{X} contains a factor-system if the following collection of subcomplexes has finite multiplicity: the smallest collection of convex subcomplexes that contains all combinatorial hyperplanes and is closed under collecting images of closest-point projection maps between its elements.) The class of cube complexes that are HHS in this way contains all universal covers of special cube complexes with finitely many immersed hyperplanes, but the cube complexes containing factor systems have not been completely characterized. In the forthcoming paper [DHS17], we show that the above question is closely related to a conjecture of the first two authors on the simplicial boundary of cube complexes [BH16, Conjecture 2.8].

More generally, we ask the following:

Question B (Factor systems in median spaces). *Is there a theory of factor systems in median spaces generalizing that in $CAT(0)$ cube complexes, such that median spaces/groups admitting factor systems are hierarchically hyperbolic?*

Presumably, a positive answer to Question B would involve the measured wallspace structure on median spaces discussed in [CDH10]. One would have to develop an analogue of the contact graph of a cube complex to serve as the underlying hyperbolic space. One must be careful since, e.g., the Baumslag-Solitar group $BS(1, 2)$ acts properly on a median space but has exponential Dehn function [Ger92] and is thus not a hierarchically hyperbolic space, by Corollary 7.5. On the other hand, if the answer to Question B is positive, one might try to do the same thing for coarse median spaces.

There are a number of other groups and spaces where it is natural to inquire whether or not they are hierarchically hyperbolic. For example:

Question C (Handlebody group). *Let H be a compact oriented 3-dimensional genus g handlebody, and let $G_g \leq MCG(\partial H)$ be the group of isotopy classes of diffeomorphisms of H . Is G_g a hierarchically hyperbolic group?*

Question D (Graph products). *Let G be a (finite) graph product of hierarchically hyperbolic groups. Is G hierarchically hyperbolic?*

The answer to Question C is presumably no, while the answer to D is most likely yes. The positive answer to Question D would follow from a strengthened version of Theorem 8.6.

There are other candidate examples of hierarchically hyperbolic spaces. For example, it is natural to ask whether a right-angled Artin group with the syllable-length metric, introduced in [KK14], which is analogous to Teichmüller space with the Weil-Petersson metric, is hierarchically hyperbolic.

As far as the difference between hierarchically hyperbolic *spaces* and *groups* is concerned, we conjecture that the following question has a positive answer:

Question E. *Is it true that the fundamental group G of a non-geometric graph manifold is a hierarchically hyperbolic group if and only if G is virtually compact special?*

It is known that G as above is virtually compact special if and only if it is chargeless in the sense of [BS05], see [HP15].

²The first question was partially answered positively in [HS16] after this paper was first posted.

There remain a number of open questions about the geometry of hierarchically hyperbolic spaces in general. Theorem 7.3 ensures, via work of Bowditch, that every asymptotic cone of a hierarchically hyperbolic space is a median space [Bow13]; further properties in this direction are established in Section 6. Motivated by combining the main result of [Sis11] on 3-manifold groups with Theorem 10.1, we ask:

Question F. *Are any two asymptotic cones of a given hierarchically hyperbolic space bi-Lipschitz equivalent?*

The notion of hierarchical quasiconvexity of a subgroup of a hierarchically hyperbolic group (G, \mathfrak{S}) generalizes quasiconvexity in word-hyperbolic groups and cubical convex-cocompactness in groups acting geometrically on CAT(0) cube complexes with factor-systems. Another notion of quasiconvexity is *stability*, defined by Durham-Taylor in [DT15]. This is a quite different notion of quasiconvexity, since stable subgroups are necessarily hyperbolic. In [DT15], the authors characterize stable subgroups of the mapping class group; it is reasonable to ask for a generalization of their results to hierarchically hyperbolic groups.

Many hierarchically hyperbolic spaces admit multiple hierarchically hyperbolic structures. However, as discussed in [BHS17a], a CAT(0) cube complex with a factor-system has a “minimal” factor-system, i.e., one that is contained in all other factor systems. In this direction, it is natural to ask whether a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ admits a hierarchically hyperbolic structure that is canonical in some way.

Recent developments. Since we posted the first version of this paper, there has been further progress on the theory of HHS and its applications.

More examples of HHS/HHG are now available, including a large class of CAT(0) cubical groups [HS16], “small-cancellation” quotients of HHGs [BHS17b], and separating curve graphs of surfaces [Vok17]. It was also recently shown by Spriano that hyperbolic spaces/groups admit alternate HHS structures that can be constructed from an arbitrary fixed collection of quasiconvex subspaces/subgroups [Spr17]. Spriano has proven additional results on modifying hierarchically hyperbolic structures to include pre-specified subgroups, under natural conditions [Spr18]. This allows him to prove that a large class of graphs of hierarchically hyperbolic groups are hierarchically hyperbolic. In the latter vein, Berlai-Robbio have, in forthcoming work [BR18], generalised the combination theorem (Theorem 8.6) in this paper, and used this to show that the class of hierarchically hyperbolic groups is closed under taking graph products.

Further developments of the theory include finiteness of the asymptotic dimension (including a quadratic upper bound for mapping class groups) [BHS17b]; a theory of boundaries generalizing the Gromov boundary of hyperbolic groups [DHS17, Mou18, Mou17]; proof of the existence of largest acylindrical actions [ABD]; and a theorem controlling quasiflats (new in both mapping class groups and cubical groups) with many applications including, for instance, a new proof of quasi-isometric rigidity for mapping class groups [BHS17c]. Mousley-Russell have recently studied Morse boundaries of hierarchically hyperbolic groups [MR18], and Abbott and the first-named author have established a linear bound on conjugator lengths in hierarchically hyperbolic groups [AB].

We stress that the present paper is foundational for almost all of the above developments; the results here are used as tools there.

Organization of the paper. Section 1 contains the full definition of a hierarchically hyperbolic space (and, more generally, a hierarchical space) and some discussion of background. Section 2 contains various basic consequences of the definition, and some tricks that are used repeatedly. In Section 3, we prove the realization theorem (Theorem 3.1). In Section 4 we establish the existence of hierarchy paths (Theorem 4.4) and the distance formula

(Theorem 4.5). Section 5 is devoted to hierarchical quasiconvexity and product regions, and Section 6 to coarse convex hulls and relatively hierarchically hyperbolic spaces. The coarse median property and its consequences are detailed in Section 7. The combination theorems for trees of spaces, graphs of groups, and products are proved in Section 8, and groups hyperbolic relative to HHG are studied in Section 9. This is applied to 3-manifolds in Section 10. Finally, in Section 11, we prove that mapping class groups are hierarchically hyperbolic.

Acknowledgments. MFH and AS thank the organizers of the Ventotene 2015 conference, and JB and MFH thank CIRM and the organizers of GAGTA 2015, where some of the work on this paper was completed. We thank Federico Berlai, Bruno Robbio, Jacob Russell and Davide Spriano for numerous comments/corrections, and anonymous referees for further helpful comments.

1. THE MAIN DEFINITION AND BACKGROUND ON HIERARCHICALLY HYPERBOLIC SPACES

1.1. **The axioms.** We begin by defining a hierarchically hyperbolic space. We will work in the context of a q -quasigeodesic space, \mathcal{X} , i.e., a metric space where any two points can be connected by a (q, q) -quasigeodesic. Obviously, if \mathcal{X} is a geodesic space, then it is a quasigeodesic space. Most of the examples we are interested in are geodesic spaces, but in order to construct hierarchically hyperbolic structures on naturally-occurring subspaces of hierarchically hyperbolic spaces, we must work in the slightly more general setting of quasigeodesic spaces.

Definition 1.1 (Hierarchically hyperbolic space). The q -quasigeodesic space $(\mathcal{X}, d_{\mathcal{X}})$ is a *hierarchically hyperbolic space* if there exists $\delta \geq 0$, an index set \mathfrak{S} , and a set $\{\mathcal{CW} : W \in \mathfrak{S}\}$ of δ -hyperbolic spaces (\mathcal{CW}, d_W) , such that the following conditions are satisfied:

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

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

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

For $V, W \in \mathfrak{S}$ satisfying $V \sqsubseteq W$ and for all $x \in \mathcal{X}$, we have:

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

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

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

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

We say that the q -quasigeodesic metric spaces $\{\mathcal{X}_i\}$ are *uniformly hierarchically hyperbolic* if each \mathcal{X}_i satisfies the axioms above and all constants, including the complexities, can be chosen uniformly. We often refer to \mathfrak{S} , together with the nesting and orthogonality relations, and the projections as a *hierarchically hyperbolic structure* for the space \mathcal{X} . Observe that \mathcal{X} is hierarchically hyperbolic with respect to $\mathfrak{S} = \emptyset$, i.e., hierarchically hyperbolic of complexity 0, if and only if \mathcal{X} is bounded. Similarly, \mathcal{X} is hierarchically hyperbolic of complexity 1 with respect to $\mathfrak{S} = \{\mathcal{X}\}$, if and only if \mathcal{X} is hyperbolic.

Notation 1.2. Where it will not cause confusion, given $U \in \mathfrak{S}$, we will often suppress the projection map π_U when writing distances in $\mathcal{C}U$, i.e., given $x, y \in \mathcal{X}$ and $p \in \mathcal{C}U$ we write $d_U(x, y)$ for $d_U(\pi_U(x), \pi_U(y))$ and $d_U(x, p)$ for $d_U(\pi_U(x), p)$. Note that when we measure distance between a pair of sets (typically both of bounded diameter) we are taking the minimum distance between the two sets. Given $A \subset \mathcal{X}$ and $U \in \mathfrak{S}$ we let $\pi_U(A)$ denote $\cup_{a \in A} \pi_U(a)$.

Remark 1.3 (Surjectivity of projections). In all of the motivating examples, and in most applications, the maps π_U are uniformly coarsely surjective.

One can always replace each $\mathcal{C}U$ with a thickening of $\pi_U(\mathcal{X})$, and hence make each π_U coarsely surjective. This is first discussed in [DHS17], where this procedure gets used; the resulting spaces are termed *normalized* hierarchically hyperbolic spaces.

More precisely, since each $\pi_U(\mathcal{X})$ is K -quasiconvex, the subset $\mathcal{C}U_{norm}$ of $\mathcal{C}U$ consisting of all geodesics that start and end in $\pi_U(\mathcal{X})$ is uniformly quasiconvex, is a (uniformly) hyperbolic geodesic metric space, and uniformly coarsely coincides with $\pi_U(\mathcal{X})$. (This “quasiconvex hull” procedure is discussed in more detail in Section 6.) Hence we can endow \mathcal{X} with a slightly different, normalized, hierarchically hyperbolic structure. Indeed, the index set is still \mathfrak{S} , each $\mathcal{C}U$ is replaced by $\mathcal{C}U_{norm}$, and the maps π_U remain unchanged (but are now coarsely surjective). Given $U, V \in \mathfrak{S}$ such that ρ_V^U is defined, we replace ρ_V^U (viewed

as a coarse map $\mathcal{CU} \rightarrow \mathcal{CV}$) with the composition $p_V \circ \rho_V^U$, where $p_V: \mathcal{CV} \rightarrow \mathcal{CU}_{norm}$ is the coarse closest-point projection.

Remark 1.4 (Surjectivity/quasiconvexity of projections in the extant applications). In the motivating examples (mapping class groups, Teichmüller space, virtually special groups, hyperbolic spaces, etc.), the projections π_U are uniformly coarsely surjective, but it is convenient to relax that requirement. As is evident from Theorem 3.1 and the key Lemma 2.6, the appropriate relaxation of coarse surjectivity is the requirement, from Definition 1.1.(1), that each $\pi_U(\mathcal{X})$ be uniformly quasiconvex in \mathcal{CU} .

In a few other places in the literature, this is not spelled out, but in each case where an issue arises, it does not affect the arguments in question. In the interest of clarity, we now summarise this as follows:

- In [DHS17, p. 4, p. 19], the authors establish a standing assumption that they are working with normalized HHSs — each π_U is uniformly coarsely surjective. In view of Remark 1.3 (or [DHS17, Proposition 1.16]), the results about normalized HHSs can be promoted to corresponding statements about general HHSs.
- In [BHS17b], Remark 1.3 allows one to assume that the HHSs in question are normalised. However, there are three places where a new HHS is constructed from an old one, and one must observe that in each of these cases, the new projections have quasiconvex image. In [BHS17b, Proposition 2.4], this holds just because the projections used in the new HHS structure coincide with those used in the old HHS structure, so quasiconvexity persists. In Proposition 6.14 and Theorem 6.2 of [BHS17b], the projections in the new HHS structures are of two types: they either coincide with projections from the old HHS structures, and thus have quasiconvex images, or they are surjective by construction.

Remark 1.5 (Large link function). It appears as though there is no actual need to require in Definition 1.1.(6) that N depend linearly on $d_W(x, x')$. Instead, we could have hypothesized that for any $C \geq 0$, there exists $N(C)$ so that the statement of the axiom holds with $N = N(C)$ whenever $d_W(x, x') \leq C$. However, one could deduce from this and the rest of the axioms that $N(C)$ grows linearly in C , so we have elected to simply build linearity into the definition.

Remark 1.6 (Summary of constants). Each hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ is associated with a collection of constants often, as above, denoted $\delta, \xi, n, \kappa_0, E, \theta_u, K$, where:

- (1) \mathcal{CU} is δ -hyperbolic for each $U \in \mathfrak{S}$,
- (2) each π_U has image of diameter at most ξ and each π_U is (K, K) -coarsely Lipschitz, $\pi_U(\mathcal{X})$ is K -quasiconvex in \mathcal{CU} , and each ρ_V^U has (image of) diameter at most ξ ,
- (3) for each $x \in \mathcal{X}$, the tuple $(\pi_U(x))_{U \in \mathfrak{S}}$ is κ_0 -consistent,
- (4) E is the larger of the constants from the bounded geodesic image axiom and the large link axiom.

Whenever working in a fixed hierarchically hyperbolic space, we use the above notation freely. We can, and shall, assume that $E \geq q, E \geq \delta, E \geq \xi, E \geq \kappa_0, E \geq K$, and $E \geq \alpha$.

Remark 1.7. We note that in Definition 1.1.(1), the assumption that the projections are Lipschitz can be replaced by the weaker assumption that there is a proper function of the projected distance which is a lower bound for the distance in the space \mathcal{X} . From this weaker assumption, the fact that the projections are actually coarsely Lipschitz then follows from the fact that we assume \mathcal{X} to be quasi-geodesic. Since the Lipschitz hypothesis is cleaner to state and, in practice, fairly easy to verify, we just remark on this for those that might find this fact useful in proving that more exotic spaces are hierarchically hyperbolic.

1.2. Comparison to the definition in [BHS17a]. Definition 1.1 is very similar to the definition of a hierarchically hyperbolic space given in [BHS17a], with the following differences:

- (1) The existence of *hierarchy paths* and the *distance formula* were stated as axioms in [BHS17a]; below, we deduce them from the other axioms. Similarly, the below *realization theorem* was formerly an axiom, but has been replaced by the (weaker) partial realization axiom.
- (2) We now require \mathcal{X} to be a quasigeodesic space. In [BHS17a], this follows from the existence of hierarchy paths, which was an axiom there.
- (3) We now require the projections $\pi_U : \mathcal{X} \rightarrow \mathcal{C}U$ to be coarsely Lipschitz; although this requirement was not imposed explicitly in [BHS17a], it follows from the distance formula, which was an axiom there.
- (4) In [BHS17a], there were five consistency inequalities; in Definition 1.1.(4), there are two. The last three inequalities in the definition from [BHS17a] follow from Proposition 1.8 below. (Essentially, the partial realization axiom has replaced part of the old consistency axiom.)
- (5) In Definition 1.1.(4), we require that, if $U \sqsubseteq V$, then $\mathbf{d}_W(\rho_W^U, \rho_W^V) \leq \kappa_0$ whenever $W \in \mathfrak{S}$ satisfies either $V \sqsubset W$ or $V \pitchfork W$ and $W \perp U$. In the context of [BHS17a], this follows by considering the standard product regions constructed using realization (see [BHS17a, Section 13.1] and Section 5.2 of the present paper).

Proposition 1.8 (ρ -consistency). *There exists κ_1 so that the following holds. Suppose that $U, V, W \in \mathfrak{S}$ satisfy both of the following conditions: $U \sqsubset V$ or $U \pitchfork V$; and $U \sqsubset W$ or $U \pitchfork W$. Then, if $V \pitchfork W$, then*

$$\min \{ \mathbf{d}_W(\rho_W^U, \rho_W^V), \mathbf{d}_V(\rho_V^U, \rho_V^W) \} \leq \kappa_1$$

and if $V \sqsubset W$, then

$$\min \{ \mathbf{d}_W(\rho_W^U, \rho_W^V), \text{diam}_{\mathcal{C}V}(\rho_V^U \cup \rho_V^W(\rho_W^U)) \} \leq \kappa_1.$$

Proof. Suppose that $U \sqsubset V$ or $U \pitchfork V$ and the same holds for U, W . Suppose that $V \pitchfork W$ or $V \sqsubseteq W$. Choose $p \in \pi_U(\mathcal{X})$. There is a uniform α so that partial realization (Definition 1.1.(8)) provides $x \in \mathcal{X}$ so that $\mathbf{d}_U(x, p) \leq \alpha$ and $\mathbf{d}_T(x, \rho_T^U) \leq \alpha$ whenever ρ_T^U is defined and coarsely constant. In particular, $\mathbf{d}_V(x, \rho_V^U) \leq \alpha$ and $\mathbf{d}_W(x, \rho_W^U) \leq \alpha$. The claim now follows from Definition 1.1.(4), with $\kappa_1 = \kappa_0 + \alpha$. \square

In view of the discussion above, we have:

Proposition 1.9. *The pair $(\mathcal{X}, \mathfrak{S})$ satisfies Definition 1.1 if and only if it is hierarchically hyperbolic in the sense of [BHS17a].*

In particular, as observed in [BHS17a]:

Proposition 1.10. *If $(\mathcal{X}, \mathfrak{S})$ is a hierarchically hyperbolic space, and \mathcal{X}' is a quasigeodesic space quasi-isometric to \mathcal{X} , then there is a hierarchically hyperbolic space $(\mathcal{X}', \mathfrak{S})$.*

1.3. A variant on the axioms. Here we introduce two slightly simpler versions of the HHS axioms and show that in the case, as in most situations which arise naturally, that the projections are coarsely surjective, then it suffices to verify the simpler axioms.

The following is a subset of the nesting axiom; here we remove the definition of the projection map $\rho_V^W : \mathcal{C}W \rightarrow 2^{\mathcal{C}V}$ in the case $V \sqsubset W$.

Definition 1.1.(2)' (Nesting variant). \mathfrak{S} is equipped with a partial order \sqsubseteq , and either $\mathfrak{S} = \emptyset$ or \mathfrak{S} contains a unique \sqsubseteq -maximal element; when $V \sqsubseteq W$, we say V is *nested* in W . We require that $W \sqsubseteq W$ for all $W \in \mathfrak{S}$. For each $W \in \mathfrak{S}$, we denote by \mathfrak{S}_W the set of $V \in \mathfrak{S}$ such that $V \sqsubseteq W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \sqsubset W$ there is a specified subset $\rho_W^V \subset \mathcal{C}W$ with $\text{diam}_{\mathcal{C}W}(\rho_W^V) \leq \xi$.

The following is a subset of the transversality and consistency axiom.

Definition 1.1.(4)' (Transversality). If $V, W \in \mathfrak{S}$ are not orthogonal and neither is nested in the other, then we say V, W are *transverse*, denoted $V \pitchfork W$. There exists $\kappa_0 \geq 0$ such that if $V \pitchfork W$, then there are sets $\rho_W^V \subseteq CW$ and $\rho_V^W \subseteq CV$ each of diameter at most ξ and satisfying:

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

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

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

The following is a variant of the bounded geodesic image axiom:

Definition 1.1.(7)' (Bounded geodesic image variant). Suppose that $x, y \in X$ and $V \sqsubset W$ have the property that there exists a geodesic from $\pi_W(x)$ to $\pi_W(y)$ which stays $(E + 2\delta)$ -far from ρ_W^V . Then $d_V(x, y) \leq E$.

Proposition 1.11. *Given a quasigeodesic space \mathcal{X} and an index set \mathfrak{S} , then $(\mathcal{X}, \mathfrak{S})$ is an HHS if it satisfies the axioms of Definition 1.1 with the following changes:*

- Replace Definition 1.1.(2) by Definition 1.1.(2)'.
- Replace Definition 1.1.(4) by Definition 1.1.(4)'.
- Replace Definition 1.1.(7) by Definition 1.1.(7)'.
- Assume that for each CU the map π_U is uniformly coarsely surjective.

Proof. To verify Definition 1.1.(2), for each $V, W \in \mathfrak{S}$ with $V \sqsubset W$, we define a map $\rho_V^W : CW \rightarrow 2^{CV}$ as follows. If $p \in CW - \mathcal{N}_E(\rho_W^V)$, then let $\rho_V^W(p) = \pi_V(x)$ for some $x \in \mathcal{X}$ with $\pi_W(x)$ (uniformly) coarsely coinciding with p . Since p does not lie E -close to ρ_W^V , this definition is coarsely independent of x by Definition 1.1.(7)'. On $\mathcal{N}_E(\rho_W^V)$, we define ρ_V^W arbitrarily. By definition, the resulting map satisfies Definition 1.1.(4). Moreover, coarse surjectivity of π_W and Definition 1.1.(7)' ensure that Definition 1.1.(7) holds. The rest of the axioms hold by hypothesis. \square

Remark 1.12. The definition of an HHS provided by Proposition 1.11 is convenient because it does not require one to define certain maps between hyperbolic spaces: Definition 1.1.(2)' is strictly weaker than Definition 1.1.(2). On the other hand, it is often convenient to work with HHS in which some of the projections π_U are not coarsely surjective; for example, this simplifies the proof that hierarchically quasiconvex subspaces inherit HHS structures in Proposition 5.6. Hence we have included both definitions.

In practice, we almost always apply consistency and bounded geodesic image in concert, which involves applying bounded geodesic image to geodesics of CW joining points in $\pi_W(\mathcal{X})$. Accordingly, Definition 1.1.(7)' is motivated by the following easy observation:

Proposition 1.13. *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. Then the conclusion of Definition 1.1.(7)' holds for all $x, y \in \mathcal{X}$ and $V, W \in \mathfrak{S}$ with $V \sqsubset W$.*

1.4. Hierarchical spaces. Although most of our focus in this paper is on hierarchically hyperbolic spaces, there are important contexts in which hyperbolicity of the spaces $CU, U \in \mathfrak{S}$ is not used; notably, this is the case for the realization theorem (Theorem 3.1). Because of the utility of a more general definition in later applications, we now define the following more general notion of a *hierarchical space*; the reader interested only in the applications to the mapping class group, 3-manifolds, cube complexes, etc. may safely ignore this subsection.

Definition 1.14 (Hierarchical space). A *hierarchical space* is a pair $(\mathcal{X}, \mathfrak{S})$ as in Definition 1.1, with \mathcal{X} a quasigeodesic space and \mathfrak{S} an index set, where to each $U \in \mathfrak{S}$ we

associate a geodesic metric space \mathcal{CU} , which we *do not require to be hyperbolic*. As before, there are coarsely Lipschitz projections $\pi_U: \mathcal{X} \rightarrow \mathcal{CU}$ and relative projections $\rho_V^U: \mathcal{CU} \rightarrow \mathcal{CV}$ whenever U, V are non-orthogonal. We require all statements in the Definition 1.1 to hold, except for hyperbolicity of the \mathcal{CU} .

Remark 1.15. Let \mathcal{X} be a quasigeodesic space that is hyperbolic relative to a collection \mathcal{P} of subspaces. Then \mathcal{X} has a hierarchical space structure: the associated spaces onto which we project are the various \mathcal{P} , together with the space $\hat{\mathcal{X}}$ obtained by coning off the elements of \mathcal{P} in \mathcal{X} . When the elements of \mathcal{P} are themselves hierarchically hyperbolic, we obtain a hierarchically hyperbolic structure on \mathcal{X} (see Section 9). Otherwise, the hierarchical structure need not be hierarchically hyperbolic since $\hat{\mathcal{X}}$ is the only one of the elements of \mathfrak{S} known to be hyperbolic.

Remark 1.16. Other than hierarchically hyperbolic spaces, we are mainly interested in hierarchical spaces $(\mathcal{X}, \mathfrak{S})$ where for all $U \in \mathfrak{S}$, except possibly when U is \sqsubseteq -minimal, we have that \mathcal{CU} is hyperbolic. This is the case, for example, in relatively hyperbolic spaces.

1.5. Consistency and partial realization points. The following definitions, which abstract the consistency inequalities from Definition 1.1.(4) and the partial realization axiom, Definition 1.1.(8), play important roles throughout our discussion. We will consider this topic in depth in Section 3.

Definition 1.17 (Consistent). Fix $\kappa \geq 0$ and let $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{\mathcal{CU}}$ be a tuple such that for each $U \in \mathfrak{S}$, the coordinate b_U is a subset of \mathcal{CU} with $\text{diam}_{\mathcal{CU}}(b_U) \leq \kappa$. The tuple \vec{b} is κ -admissible if $d_U(b_U, \pi_U(\mathcal{X})) \leq \kappa$ for all $U \in \mathfrak{S}$. The κ -admissible tuple \vec{b} is κ -consistent if, whenever $V \triangleleft W$,

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

and whenever $V \sqsubseteq W$,

$$\min \{d_W(b_W, \rho_W^V), \text{diam}_{\mathcal{CV}}(b_V \cup \rho_V^W(b_W))\} \leq \kappa.$$

In typical situations, where the maps π_U are uniformly coarsely surjective, up to a uniform enlargement of E , all tuples are admissible, so verifying consistency amounts to verifying the second condition.

Definition 1.18 (Partial realization point). Given $\theta \geq 0$ and a κ -consistent tuple \vec{b} , we say that $x \in \mathcal{X}$ is a θ -partial realization point for $\{V_j\} \subseteq \mathfrak{S}$ if

- (1) $d_{V_j}(x, b_{V_j}) \leq \theta$ for all j ,
- (2) for all j , we have $d_V(x, \rho_V^{V_j}) \leq \theta$ for any $V \in \mathfrak{S}$ with $V_j \sqsubseteq V$, and
- (3) for all W such that $W \triangleleft V_j$ for some j , we have $d_W(x, \rho_W^{V_j}) \leq \theta$.

Observe that if \vec{b} is consistent and $\{V_j\}$ is a set of pairwise-orthogonal elements, then partial realisation (Definition 1.1.(8)) provides a partial realisation point, because of admissibility.

1.6. Level. The following definition is very useful for proving statements about hierarchically hyperbolic spaces inductively. Although it is natural, and sometimes useful, to induct on complexity, it is often better to induct on level:

Definition 1.19 (Level). Let (X, \mathfrak{S}) be hierarchically hyperbolic. The level ℓ_U of $U \in \mathfrak{S}$ is defined inductively as follows. If U is \sqsubseteq -minimal then we say that its level is 1. The element U has level $k + 1$ if k is the maximal integer such that there exists $V \sqsubseteq U$ with $\ell_V = k$ and $V \neq U$. Given $U \in \mathfrak{S}$, for each $\ell \geq 0$, let \mathfrak{S}_U^ℓ be the set of $V \sqsubseteq U$ with $\ell_U - \ell_V \leq \ell$ and let $\mathfrak{T}_U^\ell = \mathfrak{S}_U^\ell - \mathfrak{S}_U^{\ell-1}$.

1.7. Maps between hierarchically hyperbolic spaces.

Definition 1.20 (Hieromorphism). Let $(\mathcal{X}, \mathfrak{S})$ and $(\mathcal{X}', \mathfrak{S}')$ be hierarchically hyperbolic structures on the spaces $\mathcal{X}, \mathcal{X}'$ respectively. A *hieromorphism*, consists of a map $f: \mathcal{X} \rightarrow \mathcal{X}'$, an injective map $f^\diamond: \mathfrak{S} \rightarrow \mathfrak{S}'$ preserving nesting, transversality, and orthogonality, and, for each $U \in \mathfrak{S}$ a map $f^*(U): \mathcal{C}U \rightarrow \mathcal{C}(f^\diamond(U))$ which is a quasi-isometric embedding where the constants are uniform over all elements of \mathfrak{S} and for which the following two diagrams coarsely commute (with uniform constants) for all nonorthogonal $U, V \in \mathfrak{S}$:

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

and

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

where $\rho_V^U: \mathcal{C}U \rightarrow \mathcal{C}V$ is the projection from Definition 1.1, which, by construction, is coarsely constant if $U \pitchfork V$ or $U \sqsubseteq V$. As the functions $f, f^*(U)$, and f^\diamond all have distinct domains, it is often clear from the context which is the relevant map; in that case we periodically abuse notation slightly by dropping the superscripts and just calling all of the maps f .

Definition 1.21 (Automorphism, hierarchically hyperbolic group). An *automorphism* of the hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ is a hieromorphism $f: (\mathcal{X}, \mathfrak{S}) \rightarrow (\mathcal{X}, \mathfrak{S})$ such that f^\diamond is bijective and each $f^*(U)$ is an isometry; hence $f: \mathcal{X} \rightarrow \mathcal{X}$ is a uniform quasi-isometry by the distance formula (Theorem 4.5).

Note that the composition of two automorphisms is again an automorphism. We say that the automorphisms f, f' are *equivalent* if $f^\diamond = (f')^\diamond$ and $f^*(U) = (f')^*(U)$ for each $U \in \mathfrak{S}$. In particular, equivalent automorphisms give equivalent quasi-isometries. Given an automorphism f , any quasi-inverse \bar{f} of f is an automorphism with $\bar{f}^\diamond = (f^\diamond)^{-1}$ and each $\bar{f}^*(U) = f^*(U)^{-1}$. Hence the set of equivalence classes of automorphisms forms a group, the *full automorphism group* of $(\mathcal{X}, \mathfrak{S})$, denoted $\text{Aut}(\mathfrak{S})$.

The finitely generated group G is *hierarchically hyperbolic* if there exists a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ and an action $G \rightarrow \text{Aut}(\mathfrak{S})$ so that the uniform quasi-action of G on \mathcal{X} is metrically proper and cobounded and \mathfrak{S} contains finitely many G -orbits. Note that if G is hierarchically hyperbolic by virtue of its action on the hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$, then (G, \mathfrak{S}) is a hierarchically hyperbolic structure with respect to any word-metric on G ; for any $U \in \mathfrak{S}$ the projection is the composition of the projection $\mathcal{X} \rightarrow \mathcal{C}U$ with a G -equivariant quasi-isometry $G \rightarrow \mathcal{X}$. In this case, (G, \mathfrak{S}) (with the implicit hyperbolic spaces and projections) is a *hierarchically hyperbolic group structure*.

Definition 1.22 (Equivariant hieromorphism). Let $(\mathcal{X}, \mathfrak{S})$ and $(\mathcal{X}', \mathfrak{S}')$ be hierarchically hyperbolic spaces and consider actions $G \rightarrow \text{Aut}(\mathfrak{S})$ and $G' \rightarrow \text{Aut}(\mathfrak{S}')$. For each $g \in G$, let $(f_g, f_g^\diamond, \{f_g^*(U)\})$ denote its image in $\text{Aut}(\mathfrak{S})$ (resp., for $g' \in G'$ we obtain $(f_{g'}, f_{g'}^\diamond, \{f_{g'}^*(U)\})$ in $\text{Aut}(\mathfrak{S}')$). Let $\phi: G \rightarrow G'$ be a homomorphism. The hieromorphism $(f, f^\diamond, \{f^*(U)\}): (\mathcal{X}, \mathfrak{S}) \rightarrow (\mathcal{X}', \mathfrak{S}')$ is ϕ -*equivariant* if for all $g \in G$ and $U \in \mathfrak{S}$, we have $f^\diamond(f_g^\diamond(U)) = f_{\phi(g)}^\diamond(f^\diamond(U))$ and the following diagram (uniformly) coarsely commutes:

$$\begin{array}{ccc}
\mathcal{C}U & \xrightarrow{f^*(U)} & \mathcal{C}(f^\diamond(U)) \\
f_g^*(U) \downarrow & & \downarrow f_{\phi(g)}^*(U) \\
\mathcal{C}(f_g^\diamond(U)) & \xrightarrow{f^*(f_g^\diamond(U))} & \mathcal{C}(f^\diamond(f_g^\diamond(U)))
\end{array}$$

In this case, $f: \mathcal{X} \rightarrow \mathcal{X}'$ is (uniformly) coarsely ϕ -equivariant in the usual sense. Also, we note for the reader that $f_g^\diamond: \mathfrak{S} \supseteq$, while $f_{\phi(g)}^\diamond: \mathfrak{S}' \supseteq$, and $f^\diamond: \mathfrak{S} \rightarrow \mathfrak{S}'$.

2. TOOLS FOR STUDYING HIERARCHICALLY HYPERBOLIC SPACES

We now collect some basic consequences of the axioms that are used repeatedly throughout the paper. However, this section need not all be read in advance. Indeed, the reader should feel free to skip this section on a first reading and return to it later when necessary. Throughout this section, we work in a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$.

2.1. Handy basic consequences of the axioms.

Lemma 2.1 (“Finite dimension”). *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space of complexity n and let $U_1, \dots, U_k \in \mathfrak{S}$ be pairwise-orthogonal. Then $k \leq n$.*

Proof. By Definition 1.1.(3), there exists $W_1 \in \mathfrak{S}$, not \sqsubseteq -maximal, so that $U_2, \dots, U_k \sqsubseteq W_1$. Applying Definition 1.1.(3) inductively yields a sequence $W_{k-1} \sqsubseteq W_{k-2} \sqsubseteq \dots \sqsubseteq W_1 \sqsubseteq S$ of distinct elements, where S is \sqsubseteq -maximal, so that $U_{i-1}, \dots, U_k \sqsubseteq W_i$ for $1 \leq i \leq k-1$. Hence $k \leq n$ by Definition 1.1.(5). \square

Lemma 2.2. *There exists χ so that $|\mathfrak{S}'| \leq \chi$ whenever $\mathfrak{S}' \subseteq \mathfrak{S}$ does not contain a pair of transverse elements.*

Proof. Let $\mathfrak{S}' \subseteq \mathfrak{S}$ be a collection of pairwise non-transverse elements, and let n be large enough that any collection of pairwise orthogonal (resp. pairwise \sqsubseteq -comparable) elements of \mathfrak{S} has cardinality at most n ; the complexity provides such an n , by Definition 1.1.(5) and Lemma 2.1. By Ramsey’s theorem, there exists N so that if $|\mathfrak{S}'| > N$ then \mathfrak{S}' contains a collection of elements, of cardinality at least $n+1$, whose elements are either pairwise orthogonal or pairwise \sqsubseteq -comparable. Hence, $|\mathfrak{S}'| \leq N$. \square

Lemma 2.3 (Consistency for pairs of points). *Let $x, y \in \mathcal{X}$ and $V, W \in \mathfrak{S}$ satisfy $V \pitchfork W$ and $d_V(x, y), d_W(x, y) > 10E$. Then, up to exchanging V and W , we have $d_V(x, \rho_V^W) \leq E$ and $d_W(y, \rho_W^V) \leq E$.*

Proof. If $d_V(x, \rho_V^W) > E$, then Definition 1.1.(4) implies $d_W(x, \rho_W^V) \leq E$. Then, either $d_W(y, \rho_W^V) \leq 9E$, in which case $d_W(x, y) \leq 10E$, a contradiction, or $d_W(y, \rho_W^V) > E$, in which case consistency implies that $d_V(y, \rho_V^W) \leq E$. \square

Corollary 2.4. *For x, y, V, W as in Lemma 2.3, and any $z \in \mathcal{X}$, there exists $U \in \{V, W\}$ such that $d_U(z, \{x, y\}) \leq 10E$.*

Proof. By Lemma 2.3, we may assume that $d_V(x, \rho_V^W), d_W(y, \rho_W^V) \leq E$. Suppose that $d_W(z, \{x, y\}) > 10E$. Then $d_W(z, \rho_W^V) > 9E$, so that, by consistency, $d_V(z, \rho_V^W) \leq E$, whence $d_V(z, x) \leq 2E$. \square

The following is needed for Theorem 3.1 and in the forthcoming [DHS17].

Lemma 2.5 (Passing large projections up the \sqsubseteq -lattice). *For every $C \geq 0$ there exists N with the following property. Let $V \in \mathfrak{S}$, let $x, y \in \mathcal{X}$, and let $\{S_i\}_{i=1}^N \subseteq \mathfrak{S}_V$ be distinct and satisfy $d_{S_i}(x, y) \geq E$. Then there exists $S \in \mathfrak{S}_V$ and i so that $S_i \sqsubset S$ and $d_S(x, y) \geq C$.*

Proof. The proof is by induction on the level k of a \sqsubseteq -minimal $S \in \mathfrak{S}_V$ into which each S_i is nested. The base case $k = 1$ is empty.

Suppose that the statement holds for a given $N = N(k)$ when the level of S is at most k . Suppose further that $|\{S_i\}| \geq N(k+1)$ (where $N(k+1)$ is a constant much larger than $N(k)$ that will be determined shortly) and there exists a \sqsubseteq -minimal $S \in \mathfrak{S}_V$ of level $k+1$ into which each S_i is nested. There are two cases.

If $d_{CS}(x, y) \geq C$, then we are done. If not, then the large link axiom (Definition 1.1.(6)) yields $K = K(C)$ and T_1, \dots, T_K , each properly nested into S (and hence of level less than $k+1$), so that any given S_i is nested into some T_j . In particular, if $N(k+1) \geq KN(k)$, there exists a j so that at least $N(k)$ elements of $\{S_i\}$ are nested into T_j . By the induction hypothesis and the finite complexity axiom (Definition 1.1.(5)), we are done. \square

The next lemma is used in the proof of Proposition 4.12, on which the existence of hierarchy paths (Theorem 4.4) relies. It is used again in Section 7 to construct coarse media.

Lemma 2.6 (Centers are consistent). *There exists κ with the following property. Let $x, y, z \in \mathcal{X}$. Let $\vec{b} = (b_W)_{W \in \mathfrak{S}}$ be so that b_W is a point in CW with the property that there exists a geodesic triangle in CW with vertices in $\pi_W(x), \pi_W(y), \pi_W(z)$ each of whose sides contains a point within distance δ of b_W . Then \vec{b} is κ -consistent.*

Proof. Recall that for $w \in \{x, y, z\}$ the tuple $(\pi_V(w))_{V \in \mathfrak{S}}$ is E -consistent. Let $U, V \in \mathfrak{S}$ be transverse. Then, by E -consistency, up to exchanging U and V and substituting z for one of x, y , we have $d_V(x, \rho_V^U), d_V(y, \rho_V^U) \leq E$, so $d_V(x, y) \leq 3E$ (recall that the diameter of ρ_V^U is at most E). Since b_V lies at distance δ from the geodesic joining $\pi_V(x), \pi_V(y)$, we have $d_V(b_V, \rho_V^U) \leq 3E + \delta$, whence the lemma holds with $\kappa = 3E + \delta$.

Suppose now $U \sqsubset V$. If b_V is within distance $10E$ of ρ_V^U , then we are done. Otherwise, up to permuting x, y, z , any geodesic $[\pi_V(x), \pi_V(y)]$ is $5E$ -far from ρ_V^U . By consistency of $(\pi_W(x)), (\pi_W(y))$ and bounded geodesic image (Definition 1.1.(7)), we have $d_U(x, y) \leq 10E$, $\text{diam}_U(\rho_V^U(\pi_V(y)) \cup \pi_U(y)) \leq E$, and $\text{diam}_U(\rho_V^U(b_V \cup \pi_V(y))) \leq 10E$. The first inequality and the definition of b_U imply $d_U(b_U, y) \leq 20E$, and taking into account the other inequalities we get $\text{diam}_U(\rho_V^U(b_V) \cup b_U) \leq 100E$.

Moreover, since $\pi_W(\mathcal{X})$ is K -quasiconvex, and b_W lies δ -close to a geodesic starting and ending in $\pi_W(\mathcal{X})$, we see that b_W lies $(K + \delta)$ -close to a point in $\pi_W(\mathcal{X})$. Hence, provided our initial choice of E was sufficiently large in terms of the constants from Definition 1.1, \vec{b} is admissible. \square

2.2. Partially ordering sets of maximal relevant elements of \mathfrak{S} . In this subsection, we describe a construction used several times in this paper, including: in the proof of realization (Theorem 3.1), in the construction of hierarchy paths (Theorem 4.4), and in the proof of the distance formula (Theorem 4.5). We expect that this construction will have numerous other applications, as is the case with the corresponding partial ordering in the case of the mapping class group, see for example [BKMM12, BM11, CLM12].

Fix $x \in \mathcal{X}$ and a tuple $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{\text{im}(\pi_U)}$, where the U -coordinate b_U is a set of diameter at most some fixed $\xi \geq 0$. For example, \vec{b} could be the tuple $(\pi_U(y))$ for some $y \in \mathcal{X}$.

In the remainder of this section, we choose $\kappa \geq 0$ and require that \vec{b} is κ -consistent. (Recall that if \vec{b} is the tuple of projections of a point in \mathcal{X} , then \vec{b} is E -consistent.)

Definition 2.7 (Relevant). First, fix $\theta \geq 100 \max\{\kappa, E\}$. Then $U \in \mathfrak{S}$ is *relevant* (with respect to x, \vec{b}, θ) if $d_U(x, b_U) > \theta$. Denote by $\mathbf{Rel}(x, \vec{b}, \theta)$ the set of relevant elements.

Let $\mathbf{Rel}_{\max}(x, \vec{b}, \theta)$ be a subset of $\mathbf{Rel}(x, \vec{b}, \theta)$ whose elements are pairwise \sqsubseteq -incomparable (for example, they could all be \sqsubseteq -maximal in $\mathbf{Rel}(x, \vec{b}, \theta)$, or they could all have the same

level). Define a relation \leq on $\mathbf{Rel}_{max}(x, \vec{b}, \theta)$ as follows. Given $U, V \in \mathbf{Rel}_{max}(x, \vec{b}, \theta)$, we have $U \leq V$ if $U = V$ or if $U \triangleleft V$ and $d_U(\rho_U^V, b_U) \leq \kappa$. Figure 1 illustrates $U < V$.




FIGURE 1. Heuristic picture of $U < V$ (for \vec{b} the coordinates of $y \in \mathcal{X}$, for concreteness). The idea is that “on the way” from x to y one “first encounters” U and is forced to change the projection from $\pi_U(x)$ to $\pi_U(y) \sim \rho_U^V$. In doing so the projection to V is not affected.

Proposition 2.8. *The relation \leq is a partial order. Moreover, either U, V are \leq -comparable or $U \perp V$.*

Proof. Clearly \leq is reflexive. Antisymmetry follows from Lemma 2.9. Suppose that U, V are \leq -incomparable. If $U \perp V$, we are done, and we cannot have $U \sqsubseteq V$ or $V \sqsubseteq U$, so suppose $U \triangleleft V$. Then, by \leq -incomparability of U, V , we have $d_U(\rho_U^V, b_U) > \kappa$ and $d_V(\rho_V^U, b_V) > \kappa$, contradicting κ -consistency of \vec{b} . This proves the assertion that transverse elements of $\mathbf{Rel}_{max}(x, \vec{b}, \theta)$ are \leq -comparable. Finally, transitivity follows from Lemma 2.10. \square

Lemma 2.9. *The relation \leq is antisymmetric.*

Proof. If $U \leq V$ and $U \neq V$, then $d_U(b_U, \rho_U^V) \leq \kappa$, so $d_U(x, \rho_U^V) > \theta - \kappa \geq 99\kappa > E$. Then, $d_V(x, \rho_V^U) \leq E$, by consistency. Thus $d_V(b_V, \rho_V^U) > \kappa$, and so, by definition $V \not\leq U$. \square

Lemma 2.10. *The relation \leq is transitive.*

Proof. Suppose that $U \leq V \leq W$. If $U = V$ or $V = W$, then $U \leq W$, and by Lemma 2.9, we cannot have $U = W$ unless $U = V = W$. Hence suppose $U \triangleleft V$ and $d_U(\rho_U^V, b_U) \leq \kappa$, while $V \triangleleft W$ and $d_V(\rho_V^W, b_V) \leq \kappa$. By the definition of $\mathbf{Rel}_{max}(x, \vec{b}, \theta)$, we have $d_T(x, b_T) > 100\kappa$ for $T \in \{U, V, W\}$.

We first claim that $d_V(\rho_V^U, \rho_V^W) > 10E$. Indeed, $d_U(b_U, \rho_U^V) \leq \kappa$, so $d_U(\rho_U^V, x) \geq 90\kappa$, whence $d_V(\rho_V^U, x) \leq E \leq \kappa$ by E -consistency of the tuple $(\pi_T(x))_{T \in \mathcal{S}}$. On the other hand, $d_V(\rho_V^W, b_V) \leq \kappa$, so $d_V(\rho_V^U, \rho_V^W) > 10E$ as claimed. Hence, by Lemma 2.11, we have $U \triangleleft W$.

Since $\text{diam}(\text{im}(\pi_W)) > 100\kappa$ — indeed, $d_W(x, b_W) > 100\kappa$ and $b_W \in \text{im}(\pi_W(\mathcal{X}))$ — partial realization (Definition 1.1.(8)) provides $a \in \mathcal{X}$ satisfying $d_W(a, \{\rho_W^U, \rho_W^V\}) \geq 10\kappa$.

We thus have $d_U(a, \rho_U^W) \leq E$ by E -consistency of $(\pi_T(a))_{T \in \mathcal{S}}$, and the same is true with V replacing U . Hence $d_V(\rho_V^U, a) > E$, so consistency implies $d_U(a, \rho_U^V) \leq E$. Thus $d_U(\rho_U^V, \rho_U^W) \leq 2E$. Thus $d_U(b_U, \rho_U^W) \leq 2E + \kappa < 10\kappa$, whence $d_U(x, \rho_U^W) > 50\kappa > E$, so $d_W(x, \rho_W^U) \leq E$ by consistency and the fact that $U \triangleleft W$. It follows that $d_W(b_W, \rho_W^U) \geq 100\kappa - E > \kappa$, so, again by consistency, $d_U(b_U, \rho_U^W) \leq \kappa$, i.e., $U \leq W$. \square

Lemma 2.11. *Let $U, V, W \in \mathfrak{S}$ satisfy $\text{diam}(\text{im}(\pi_U)), \text{diam}(\text{im}(\pi_V)), \text{diam}(\text{im}(\pi_W)) > 10E$, and $U \triangleleft V, W \triangleleft V$, and $\text{d}_V(\rho_V^U, \rho_V^W) > 10E$. Suppose moreover that U and W are \sqsubseteq -incomparable. Then $U \triangleleft W$.*

Proof. If $U \perp W$, then by the partial realization axiom (Definition 1.1.(8)) and the lower bound on diameters, there exists an E -partial realization point x for $\{U, W\}$ so that

$$\text{d}_U(\rho_U^V, x), \text{d}_W(\rho_W^V, x) > E.$$

This contradicts consistency since $\text{d}_V(\rho_V^U, \rho_V^W) > 10E$; indeed, by consistency $\text{d}_V(\rho_V^U, x) \leq E, \text{d}_V(\rho_V^W, x) \leq E$, i.e., $\text{d}_V(\rho_V^U, \rho_V^W) \leq 2E$. Hence $U \triangleleft W$. \square

2.3. Coloring relevant elements. In this subsection, the key result is Lemma 2.14, which we will apply in proving the existence of hierarchy paths in Section 4.3.

Fix $x, y \in \mathcal{X}$. As above, let $\mathbf{Rel}(x, y, 100E)$ consist of those $V \in \mathfrak{S}$ for which $\text{d}_V(x, y) > 100E$. Recall that, given $U \in \mathfrak{S}$, we denote by \mathfrak{T}_U^ℓ the set of $V \in \mathfrak{S}_U$ such that $\ell_U - \ell_V = \ell$. In particular, if $V, V' \in \mathfrak{T}_U^\ell$ and $V \sqsubseteq V'$, then $V = V'$. Let $\mathbf{Rel}_U^\ell(x, y, 100E) = \mathbf{Rel}(x, y, 100E) \cap \mathfrak{T}_U^\ell$, the set of $V \sqsubseteq U$ so that $\text{d}_V(x, y) > 100E$ and $\ell_U - \ell_V = \ell$.

By Proposition 2.8, the relation \leq on $\mathbf{Rel}_U^\ell(x, y, 100E)$ defined as follows is a partial order: $V \leq V'$ if either $V = V'$ or $\text{d}_V(y, \rho_{V'}^V) \leq E$.

Definition 2.12 (Relevant graph). Denote by \mathcal{G} the graph with vertex-set $\mathbf{Rel}_U^\ell(x, y, 100E)$, with two vertices adjacent if and only if the corresponding elements of $\mathbf{Rel}_U^\ell(x, y, 100E)$ are orthogonal. Let \mathcal{G}^c denote the complementary graph of \mathcal{G} , i.e., the graph with the same vertices and edges corresponding to \leq -comparability.

The next lemma is an immediate consequence of Proposition 2.8:

Lemma 2.13. *Elements of $V, V' \in \mathbf{Rel}_U^\ell(x, y, 100E)$ are adjacent in \mathcal{G} if and only if they are \leq -incomparable.*

Lemma 2.14 (Coloring relevant elements). *Let χ be the maximal cardinality of a set of pairwise orthogonal elements of \mathfrak{T}_U^ℓ . Then there exists a χ -coloring of the set of relevant elements of \mathfrak{T}_U^ℓ such that non-transverse elements have different colors.*

Proof. Since each clique in \mathcal{G} — i.e., each \leq -antichain in $\mathbf{Rel}_U^\ell(x, y, 100E)$ — has cardinality at most χ , Dilworth's theorem [Dil50, Theorem 1.1] implies that \mathcal{G} can be colored with χ colors in such a way that \leq -incomparable elements have different colors; hence non-transverse elements have different colors. \square

Remark 2.15. The constant χ provided by Lemma 2.14 is bounded by the complexity of $(\mathcal{X}, \mathfrak{S})$, by Lemma 2.2.

3. REALIZATION OF CONSISTENT TUPLES

The goal of this section is to prove Theorem 3.1. In this section we will work with a fixed hierarchical space $(\mathcal{X}, \mathfrak{S})$. We will use the concepts of consistency and partial realization points; see Definition 1.17 and Definition 1.18.

Theorem 3.1 (Realization of consistent tuples). *For each $\kappa \geq 1$ there exist $\theta_e, \theta_u \geq 0$ such that the following holds. Let $\vec{b} \in \prod_{W \in \mathfrak{S}} 2^{CW}$ be κ -consistent; for each W , let b_W denote the CW -coordinate of \vec{b} .*

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

Remark 3.2. In typical cases, where the π_U are uniformly coarsely surjective, the admissibility part of the consistency hypothesis is satisfied automatically.

Proof of Theorem 3.1. The main task is to prove the following claim about a κ -consistent admissible tuple \vec{b} :

Claim 1. Let $\{V_j\}$ be a family of pairwise orthogonal elements of \mathfrak{S} , all of level at most ℓ . Then there exists $\theta_e = \theta_e(\ell, \kappa) > 100E\kappa\alpha$ and pairwise-orthogonal $\{U_i\}$ so that:

- (1) each U_i is nested into some V_j ,
- (2) for each V_j there exists some U_i nested into it, and
- (3) any E -partial realization point x for $\{U_i\}$ satisfies $d_W(b_W, x) \leq \theta_e$ for each $W \in \mathfrak{S}$ for which there exists j with $W \sqsubseteq V_j$.

Applying Claim 1 when $\ell = \ell_S$, where $S \in \mathfrak{S}$ is the unique \sqsubseteq -maximal element, along with the Partial Realization axiom (Definition 1.1.(8)), completes the existence proof, giving us a constant θ_e . If x, y both have the desired property, then $d_V(x, y) \leq 2\theta_e + \kappa$ for all $V \in \mathfrak{S}$, whence the uniqueness axiom (Definition 1.1.(9)) ensures that $d(x, y) \leq \theta_u$, for an appropriate θ_u . Hence to prove the theorem it remains to prove Claim 1, which we do now.

The claim when $\ell = 1$ follows from admissibility and the partial realization axiom (Definition 1.1.8), so we assume that the claim holds for $\ell - 1 \geq 1$, with $\theta_e(\ell - 1, \kappa) = \theta'_e$, and prove it for level ℓ .

Reduction to the case $|\{V_j\}| = 1$: It suffices to prove the claim in the case where $\{V_j\}$ has a single element, V . To see this, note that once we prove the claim for each V_j separately, yielding a collection of pairwise-orthogonal sets $\{U_i^j \sqsubseteq V_j\}$ with the desired properties, then we take the union of these sets to obtain the claim for the collection $\{V_j\}$.

The case $\{V_j\} = \{V\}$: Fix $V \in \mathfrak{S}$ so that $\ell_V = \ell$. If for each $x \in \mathcal{X}$ that satisfies $d_V(x, b_V) \leq E$ we have $d_W(b_W, x) \leq 100E\kappa\alpha$ for $W \in \mathfrak{S}_V$, then the claim follows with $\{U_i\} = \{V\}$. Hence, we can suppose that this is not the case.

We are ready for the main argument, which is contained in Lemma 3.3 below. We will construct $\{U_i\}$ incrementally, using Lemma 3.3, which essentially says that either we are done at a certain stage or we can add new elements to $\{U_i\}$.

We will say that the collection \mathfrak{U} of elements of \mathfrak{S}_V is *totally orthogonal* if any pair of distinct elements of \mathfrak{U} are orthogonal. Given a totally orthogonal family \mathfrak{U} we say that $W \in \mathfrak{S}_V$ is \mathfrak{U} -*generic* if there exists $U \in \mathfrak{U}$ so that W is not orthogonal to U . Notice that no W is \emptyset -generic.

A totally orthogonal collection $\mathfrak{U} \subseteq \mathfrak{S}_V$ is C -*good* if any E -partial realization point x for \mathfrak{U} has the property that for each $W \in \mathfrak{S}_V$ we have $d_W(x, b_W) \leq C$. (Notice that our goal is to find such \mathfrak{U} .) A totally orthogonal collection $\mathfrak{U} \subseteq \mathfrak{S}_V$ is C -*generically good* if any E -partial realization point x for \mathfrak{U} has the property that for each \mathfrak{U} -generic $W \in \mathfrak{S}_V$ we have $d_W(x, b_W) \leq C$ (e.g., for $\mathfrak{U} = \emptyset$).

We can now quickly finish the proof of the claim using Lemma 3.3 about extending generically good sets, which we state and prove below. Start with $\mathfrak{U} = \emptyset$. If \mathfrak{U} is C -good for $C = 100E\kappa\alpha$, then we are done. Otherwise we can apply Lemma 3.3 and get $\mathfrak{U}_1 = \mathfrak{U}'$ as in the lemma. Inductively, if \mathfrak{U}_n is not $10^n C$ -good, we can apply the lemma and extend \mathfrak{U}_n to a new totally orthogonal set \mathfrak{U}_{n+1} . Since there is a bound on the cardinality of totally orthogonal sets by Lemma 2.1, in finitely many steps we necessarily get a good totally orthogonal set, and this concludes the proof of the claim, and hence of the theorem. \square

Lemma 3.3. *For every $C \geq 100E\kappa\alpha$ the following holds. Let $\mathfrak{U} \subseteq \mathfrak{S}_V - \{V\}$ be totally orthogonal and C -generically good but not C -good. Then there exists a totally orthogonal, $10C$ -generically good collection $\mathfrak{U}' \subseteq \mathfrak{S}_V$ with $\mathfrak{U} \subsetneq \mathfrak{U}'$.*

Proof. Let x_0 be an E -partial realization point for \mathfrak{U} so that there exists some $W \sqsubseteq V$ for which $d_W(b_W, x_0) > C$.

The idea is to try to “move towards” \vec{b} starting from x_0 , by looking at all relevant elements of \mathfrak{S}_V that lie between them and finding out which ones are the “closest” to \vec{b} .

Let \mathfrak{V}_{max} be the set of all $W \sqsubseteq V$ for which:

- (1) $d_W(b_W, x_0) > C$ and
- (2) W is not properly nested into any element of \mathfrak{S}_V satisfying the above inequality.

We now establish two facts about \mathfrak{V}_{max} .

Applying Proposition 2.8 to partially order \mathfrak{V}_{max} : For $U, U' \in \mathfrak{V}_{max}$, write $U \leq U'$ if either $U = U'$ or $U \triangleleft U'$ and $d_U(\rho_U^{U'}, b_U) \leq 10E\kappa$; this is a partial order by Proposition 2.8, which also implies that if $U, U' \in \mathfrak{V}_{max}$ are transverse then they are \leq -comparable. Hence any two \leq -maximal elements of \mathfrak{V}_{max} were orthogonal, and we denote by \mathfrak{V}'_{max} the set of \leq -maximal (hence pairwise-orthogonal) elements of \mathfrak{V}_{max} .

Finiteness of \mathfrak{V}_{max} : We now show that $|\mathfrak{V}_{max}| < \infty$. By Lemma 2.2 and Ramsey’s theorem, if \mathfrak{V}_{max} was infinite then it would contain an infinite subset of pairwise transverse elements, so, in order to conclude that $|\mathfrak{V}_{max}| < \infty$, it suffices to bound the cardinality of a pairwise-transverse subset of \mathfrak{V}_{max} .

Suppose that $W_1 < \dots < W_s \in \mathfrak{V}_{max}$ are pairwise transverse. By partial realization (Definition 1.1.(8)) and admissibility, there exists $z \in \mathcal{X}$ such that $d_{W_s}(z, b_{W_s}) \leq \alpha$ and $d_{W_i}(\rho_{W_i}^{W_s}, z) \leq \alpha$ for each $i \neq s$, and such that $d_V(z, \rho_V^{W_s}) \leq \alpha$. By consistency of \vec{b} and bounded geodesic image, $\rho_V^{W_s}$ has to be within distance $10E\kappa$ of a geodesic in $\mathcal{C}V$ from x_0 to b_V . In particular $d_V(x_0, z) \leq \theta'_e + 100E\kappa\alpha + 10E\kappa$. Also, for each $i \neq s$,

$$\begin{aligned} d_{W_i}(x_0, z) &\geq d_{W_i}(x_0, b_{W_i}) - d_{W_i}(b_{W_i}, \rho_{W_i}^{W_s}) - d_{W_i}(\rho_{W_i}^{W_s}, z) \\ &\geq 100E\kappa\alpha - 10E\kappa - \alpha \geq 50E\kappa\alpha \geq 50E. \end{aligned}$$

Indeed, $d_{W_i}(b_{W_i}, \rho_{W_i}^{W_s}) \leq 10E\kappa$ since $W_i < W_s$, while $d_{W_i}(\rho_{W_i}^{W_s}, z) \leq \alpha$ by our choice of z . Lemma 2.5 now provides the required bound on s .

Choosing \mathfrak{U}' : Since $\ell_U < \ell_V$ for all $U \in \mathfrak{V}'_{max}$, by induction there exists a totally orthogonal set $\{U_i\}$ so that any E -partial realization point x for $\{U_i\}$ satisfies $d_T(b_T, x) \leq \theta'_e$ for each $T \in \mathfrak{S}$ nested into some $U \in \mathfrak{V}'_{max}$. Let $\mathfrak{U}' = \{U_i\} \cup \mathfrak{U}$.

Choose such a partial realization point x and let $W \sqsubseteq V$ be \mathfrak{U}' -generic. Our goal is to bound $d_W(x, b_W)$, and we will consider 4 cases.

If there exists $U \in \mathfrak{U}$ that is not orthogonal to W , then we are done by hypothesis, since any E -partial realization point for \mathfrak{U}' is also an E -partial realization point for \mathfrak{U} .

Hence, from now on, assume that W is orthogonal to each $U \in \mathfrak{U}$, i.e. W is not \mathfrak{U} -generic.

If $W \sqsubseteq U$ for some $U \in \mathfrak{V}'_{max}$, then we are done by induction.

Suppose that $W \triangleleft U$ for some $U \in \mathfrak{V}'_{max}$. For each $U_i \sqsubseteq U$ — and our induction hypothesis implies that there is at least one such U_i — we have $d_W(x, \rho_W^{U_i}) \leq E$ since x is a partial realization point for $\{U_i\}$ and either $U_i \sqsubseteq W$ or $U_i \triangleleft W$ (since W is \mathfrak{U}' -generic but not \mathfrak{U} -generic). The triangle inequality therefore yields:

$$d_W(x, b_W) \leq E + d_W(\rho_W^{U_i}, \rho_W^U) + d_W(b_W, \rho_W^U).$$

By Definition 1.1.(4), $d_W(\rho_W^{U_i}, \rho_W^U) \leq E$, and we will show that $d_W(b_W, \rho_W^U) \leq 2C$, so that $d_W(x, b_W) \leq 2E + 2C$.

Suppose, for a contradiction, that $d_W(b_W, \rho_W^U) > 2C$. If $d_U(\rho_U^W, x_0) \leq E$, then

$$d_U(\rho_U^W, b_U) \geq C - E > \kappa,$$

by consistency, whence $d_W(\rho_W^U, b_W) \leq \kappa$, a contradiction.

On the other hand, if $d_U(\rho_U^W, x_0) > E$, then $d_W(x_0, \rho_W^U) \leq E$ by consistency. Hence $d_W(x_0, b_W) \geq 2C - E$. Hence there exists a \sqsubseteq -maximal $W' \neq V$ with the property that $W \sqsubseteq W' \sqsubseteq V$ and $d_{W'}(x_0, b_{W'}) > C$ (possibly $W' = W$). Such a W' is in \mathfrak{W}_{max} by definition.

Since $W \triangleleft U$, and W' and U are \sqsubseteq -incomparable, $W' \triangleleft U$. Thus U and W' are \leq -comparable, by Proposition 2.8. Since $W' \neq U$ and U is \leq -maximal, we have $W' \leq U$, i.e., $d_{W'}(b_{W'}, \rho_{W'}^U) \leq 10E\kappa$. Since \leq is antisymmetric, by Lemma 2.9, we have $d_U(b_U, \rho_U^{W'}) > 10E\kappa$. Since $d_U(\rho_U^W, \rho_U^{W'}) \leq E$ (from Definition 1.1.(4)), we have $d_U(b_U, \rho_U^W) > 10E\kappa - E > \kappa$, since $E \geq 1$, so, by consistency, $d_W(b_W, \rho_W^U) \leq \kappa$, a contradiction.

Finally, suppose $U \not\sqsubseteq W$ for some $U \in \mathfrak{W}'_{max}$. In this case, by \sqsubseteq -maximality of U , we have $d_W(x_0, b_W) \leq C$. Also, $d_W(x, \rho_W^{U_i}) \leq E$ for any $U_i \sqsubseteq U$ since x is a partial realization point, so that $d_W(x, \rho_W^U) \leq 2E$, since $d_W(\rho_W^U, \rho_W^{U_i}) \leq E$ by Definition 1.1.(4). If $d_W(x, b_W) > 2C$, then we claim $d_U(x_0, b_U) \leq 10E\kappa$, a contradiction. Indeed, any geodesic in CW from $\pi_W(x_0)$ to b_W does not enter the E -neighborhood of ρ_W^U . By bounded geodesic image, $\text{diam}_U(\rho_U^W(\pi_W(x_0)) \cup \rho_U^W(b_W)) \leq E$ and by consistency, $\text{diam}_U(\rho_U^W(\pi_W(x_0)) \cup \pi_U(x_0)) \leq E$ and $\text{diam}_U(\rho_U^W(b_W) \cup b_U) \leq \kappa$, and we obtain the desired bound on $d_U(x_0, b_U)$. This completes the proof of the lemma. \square

4. HIERARCHY PATHS AND THE DISTANCE FORMULA

Throughout this section, fix a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$.

4.1. Definition of hierarchy paths and statement of main theorems. Our goal is to deduce the existence of hierarchy paths (Theorem 4.4) from the other axioms and to prove the distance formula (Theorem 4.5).

Definition 4.1 (Quasigeodesic, unparameterized quasigeodesic). A (D, D) -*quasigeodesic* in the metric space M is a (D, D) -quasi-isometric embedding $f: [0, \ell] \rightarrow M$; we allow f to be a coarse map, i.e., to send points in $[0, \ell]$ to uniformly bounded sets in M . A (coarse) map $f: [0, \ell] \rightarrow M$ is a (D, D) -*unparameterized quasigeodesic* if there exists a strictly increasing function $g: [0, L] \rightarrow [0, \ell]$ such that $g(0) = f(0), g(L) = f(\ell)$, and $f \circ g: [0, L] \rightarrow M$ is a (D, D) -quasigeodesic and for each $j \in [0, L] \cap \mathbb{N}$, we have $\text{diam}_M(f(g(j)) \cup f(g(j+1))) \leq D$.

Definition 4.2 (Hierarchy path). For $D \geq 1$, a (not necessarily continuous) path $\gamma: [0, \ell] \rightarrow \mathcal{X}$ is a D -*hierarchy path* if

- (1) γ is a (D, D) -quasigeodesic,
- (2) for each $W \in \mathfrak{S}$, the path $\pi_W \circ \gamma$ is an unparameterized (D, D) -quasigeodesic.

Notation 4.3. Given $A, B \in \mathbb{R}$, we denote by $\{A\}_B$ the quantity which is A if $A \geq B$ and 0 otherwise. Given C, D , we write $A \asymp_{C,D} B$ to mean $C^{-1}A - D \leq B \leq CA + D$.

Theorem 4.4 (Existence of Hierarchy Paths). *Let $(\mathcal{X}, \mathfrak{S})$ be hierarchically hyperbolic. Then there exists D_0 so that any $x, y \in \mathcal{X}$ are joined by a D_0 -hierarchy path.*

Theorem 4.5 (Distance Formula). *Let $(\mathcal{X}, \mathfrak{S})$ 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 $x, y \in \mathcal{X}$,*

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

The proofs of the above two theorems are intertwined, and we give the proof immediately below. This relies on several lemmas, namely Lemma 4.11, proved in Section 4.3, and Lemmas 4.19 and 4.18, proved in Section 4.4.

Proof of Theorems 4.5 and 4.4. The lower bound demanded by Theorem 4.5 is given by Lemma 4.19 below. By Lemma 4.11 and Lemma 4.18, there is a monotone path (see Definition 4.8) whose length realizes the upper bound on $d_{\mathcal{X}}(x, y)$, and the same holds for any subpath of this path, which is therefore a hierarchy path, proving Theorem 4.4 and completing the proof of Theorem 4.5. \square

4.2. Good and proper paths: definitions. We now define various types of (non-continuous) paths in \mathcal{X} that will appear on the way to hierarchy paths.

Definition 4.6 (Discrete path). A K -discrete path is a map $\gamma: I \rightarrow \mathcal{X}$, where I is an interval in \mathbb{Z} and $d_{\mathcal{X}}(\gamma(i), \gamma(i+1)) \leq K$ whenever $i, i+1 \in I$. The length $|\alpha|$ of a discrete path α is $\max I - \min I$.

Definition 4.7 (Efficient path). A discrete path α with endpoints x, y is K -efficient if $|\alpha| \leq K d_{\mathcal{X}}(x, y)$.

Definition 4.8 (Monotone path). Given $U \in \mathfrak{S}$, a K -discrete path α and a constant L , we say that α is L -monotone in U if whenever $i \leq j$ we have $d_U(\alpha(0), \alpha(i)) \leq d_U(\alpha(0), \alpha(j)) + L$. A path which is L -monotone in U for all $U \in \mathfrak{S}$ is said to be L -monotone.

Definition 4.9 (Good path). A K -discrete path that is L -monotone in U is said to be (K, L) -good for U . Given $\mathfrak{S}' \subseteq \mathfrak{S}$, a path α that is (K, L) -good for each $V \in \mathfrak{S}'$ is (K, L) -good for \mathfrak{S}' .

Definition 4.10 (Proper path). A discrete path $\alpha: \{0, \dots, n\} \rightarrow \mathcal{X}$ is (r, K) -proper if for $0 \leq i < n-1$, we have $d_{\mathcal{X}}(\alpha(i), \alpha(i+1)) \in [r, r+K]$ and $d_{\mathcal{X}}(\alpha(n-1), \alpha(n)) \leq r+K$. Observe that (r, K) -properness is preserved by passing to subpaths.

4.3. Good and proper paths: existence. Our goal in this subsection is to join points in \mathcal{X} with proper paths, i.e., to prove Lemma 4.11. This relies on the much more complicated Proposition 4.12, which produces good paths (which are then easily made proper).

Lemma 4.11. *There exists K so that for any $r \geq 0$, any $x, y \in \mathcal{X}$ are joined by a K -monotone, (r, K) -proper discrete path.*

Proof. Let $\alpha_0: \{0, \dots, n_0\} \rightarrow \mathcal{X}$ be a K -monotone, K -discrete path joining x, y , which exists by Proposition 4.12. We modify α_0 to obtain the desired path in the following way. Let $j_0 = 0$ and, proceeding inductively, let j_i be the minimal $j \leq n$ such that either $d_{\mathcal{X}}(\alpha_0(j_{i-1}), \alpha_0(j)) \in [r, r+K]$ or $j = n$. Let m be minimal so that $j_m = n$ and define $\alpha: \{0, \dots, m\} \rightarrow \mathcal{X}$ by $\alpha(j) = \alpha_0(j_j)$. The path α is (r, K) -proper by construction; it is easily checked that K -monotonicity is not affected by the above modification; the new path is again discrete, although for a larger discreteness constant. \square

It remains to establish the following proposition, whose proof is postponed until the end of this section, after several preliminary statements have been obtained.

Proposition 4.12. *There exists K so that any $x, y \in \mathcal{X}$ are joined by path that is (K, K) -good for each $U \in \mathfrak{S}$.*

Definition 4.13 (Hull of a pair of points). For each $x, y \in \mathcal{X}, \theta \geq 0$, let $H_{\theta}(x, y)$ 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 a geodesic in CW joining $\pi_W(x)$ to $\pi_W(y)$. Note that $x, y \in H_{\theta}(x, y)$.

Remark 4.14. The notion of a hull is generalized in Section 6 to hulls of arbitrary finite sets, but we require only the version for pairs of points in this section.

Lemma 4.15 (Retraction onto hulls). *There exist $\theta, K \geq 0$ such that, for each $x, y \in \mathcal{X}$, there exists a (K, K) -coarsely Lipschitz map $r : \mathcal{X} \rightarrow H_\theta(x, y)$ that restricts to the identity on $H_\theta(x, y)$.*

Proof. Let κ be the constant from Lemma 2.6, let θ_e be chosen as in the realization theorem (Theorem 3.1), and let $p \in \mathcal{X} - H_{\theta_e}(x, y)$. Define a tuple $\vec{b} = (b_W^p) \in \prod_{W \in \mathfrak{S}} 2^{CW}$ so that b_W^p is on a geodesic in CW from $\pi_W(x)$ to $\pi_W(y)$ and is within distance δ of the other two sides of a triangle with vertices in $\pi_W(x), \pi_W(y), \pi_W(p)$. By Lemma 2.6, this is a consistent tuple. Hence, by the realization theorem (Theorem 3.1), there exists $r(p) \in H_{\theta_e}(x, y)$ so that $d_W(\pi_W(r(p)), b_W^p) \leq \theta_e$. For $p \in H_{\theta_e}(x, y)$, let $r(p) = p$.

To see that r is coarsely Lipschitz, it suffices to bound $d_{\mathcal{X}}(r(p), r(q))$ when $p, q \in \mathcal{X}$ satisfy $d_{\mathcal{X}}(p, q) \leq 1$. For such p, q we have $d_W(b_W^p, b_W^q) \leq 100E$, so that Theorem 3.1 implies $d_{\mathcal{X}}(r(p), r(q)) \leq \theta_u(100E)$, as required. \square

Corollary 4.16. *There exist $\theta, K \geq 0$ such that, for each $x, y \in \mathcal{X}$, there exists a K -discrete and K -efficient path that lies in $H_\theta(x, y)$ and joins x to y .*

Proof. We can assume that $d_{\mathcal{X}}(x, y) \geq 1$. Since \mathcal{X} is a quasigeodesic space, there exists $C = C(\mathcal{X}) \geq 1$ and a (C, C) -quasi-isometric embedding $\gamma : [0, L] \rightarrow \mathcal{X}$ with $\gamma(0) = x, \gamma(L) = y$. Let ρ be the path obtained by restricting $r \circ \gamma : [0, L] \rightarrow H_\theta(x, y)$ to $[0, L] \cap \mathbb{N}$, where r is the retraction obtained in Lemma 4.15. Then $d_{\mathcal{X}}(\rho(i), \rho(i+1)) \leq 10KC$ since r is (K, K) -coarsely Lipschitz and γ is (C, C) -coarsely Lipschitz, i.e., ρ is $10KC$ -discrete. Finally, ρ is efficient because $L \leq Cd_{\mathcal{X}}(x, y) + C \leq 2Cd_{\mathcal{X}}(x, y)$. \square

The efficiency part of the corollary is used in Lemma 4.19.

4.3.1. *Producing good paths.* We will need the following lemma, which is a special case of Proposition 6.4.2. We give a proof in the interest of a self-contained exposition.

Lemma 4.17. *For any θ_0 there exists a constant θ such that for every $x, y \in \mathcal{X}$ and every $x', y' \in H_{\theta_0}(x, y)$, we have $H_{\theta_0}(x', y') \subseteq H_\theta(x, y)$.*

Proof. For any $z \in H_{\theta_0}(x', y')$ and $W \in \mathfrak{S}$ the projection $\pi_W(z)$ lies $2(\delta + \theta_0)$ -close to a geodesic in CW from $\pi_W(x)$ to $\pi_W(y)$, by a thin quadrilateral argument. \square

We now prove the main proposition of this subsection.

Proof of Proposition 4.12. Recall that, for $\ell \geq 0$ and $U \in \mathfrak{S}$, the set \mathfrak{S}_U^ℓ consists of those $V \in \mathfrak{S}_U$ with $\ell_U - \ell_V \leq \ell$, and that \mathfrak{T}_U^ℓ consists of those $V \in \mathfrak{S}_U$ with $\ell_U - \ell_V = \ell$.

We prove by induction on ℓ that there exist θ, K such that for any $\ell \geq 0, x, y \in \mathcal{X}$ and $U \in \mathfrak{S}$, there is a path α in $H_\theta(x, y)$ connecting x to y such that α is (K, K) -good for \mathfrak{S}_U^ℓ . It then follows that for any $x, y \in \mathcal{X}$, there exists a path α in $H_\theta(x, y)$ connecting x to y such that α is (K, K) -good for \mathfrak{S} ; this latter statement directly implies the proposition.

For $a, b \in \mathcal{X}$, denote by $[a, b]_W$ a geodesic in CW from $\pi_W(a)$ to $\pi_W(b)$. Fix $U \in \mathfrak{S}$.

The case $\ell = 0$: In this case, $\mathfrak{S}_U^0 = \{U\}$. By Corollary 4.16, there exists θ_0, K and a K -discrete, K -efficient path $\alpha'_0 : \{0, \dots, k\} \rightarrow H_{\theta_0}(x, y)$ joining x to y .

Similarly, for each $x', y' \in H_{\theta_0}(x, y)$ there exists a K -discrete path β contained in $H_{\theta_0}(x', y')$, joining x' to y' , and recall that $H_{\theta_0}(x', y')$ is contained in $H_\theta(x, y)$ for a suitable θ in view of Lemma 4.17.

We use the term *straight path* to refer to a path, such as β , which for each $V \in \mathfrak{S}$ projects uniformly close to a geodesic of $\mathcal{C}(V)$.

We now fix $U \in \mathfrak{S}$, and, using the observation in the last paragraph explain how to modify α'_0 to obtain a K -discrete path α_0 in $H_\theta(x, y)$ that is K -monotone in U ; the construction will rely on replacing problematic subpaths with straight paths.

A point $t \in \{0, \dots, k\}$ is a U -omen if there exists $t' > t$ so that $d_U(\alpha'_0(0), \alpha'_0(t)) > d_U(\alpha'_0(0), \alpha'_0(t')) + 5KE$. If α'_0 has no U -omens, then we can take $\alpha_0 = \alpha'_0$, so suppose that there is a U -omen and let t_0 be the minimal U -omen, and let $t'_0 > t_0$ be maximal so that $d_U(\alpha'_0(0), \alpha'_0(t_0)) > d_U(\alpha'_0(0), \alpha'_0(t'_0))$. Inductively define t_j to be the minimal U -omen with $t_j \geq t'_{j-1}$, if such t_j exists; and when t_j exists, we define t'_j to be maximal in $\{0, \dots, k\}$ satisfying $d_U(\alpha'_0(0), \alpha'_0(t_j)) > d_U(\alpha'_0(0), \alpha'_0(t'_j))$. For each $j \geq 0$, let $x'_j = \alpha'_0(t_j)$ and let $y'_j = \alpha'_0(t'_j)$. See Figure 2.




FIGURE 2. The picture above shows part of α'_0 in \mathcal{X} , and that below shows its projection to U . The point t_j is an omen, as witnessed by the point marked with a square. Inserting the dashed path β_j , and deleting the corresponding subpath of α'_0 , makes t_j cease to be an omen.

For each j , there exists a K -discrete path β_j which lies in $H_{\theta_0}(x'_j, y'_j) \subseteq H_\theta(x, y)$ and is a straight path from x'_j to y'_j . Let α_0 be obtained from α'_0 by replacing each $\alpha'_0([t_j, t'_j])$ with β_j . Clearly, α_0 connects x to y , is K -discrete, and is contained in $H_\theta(x, y)$. For each j we have that $\text{diam}_{CU}(\beta_j) \leq d_U(x'_j, y'_j) + 2\theta_0$.

Notice that $d_U(x'_j, y'_j) < 2KE + 10\theta_0$. In fact, since $\alpha'_0(0), \alpha'_0(t_j), \alpha'_0(t'_j)$ lie θ_0 -close to a common geodesic and $d_U(\alpha'_0(0), \alpha'_0(t_j)) \geq d_U(\alpha'_0(0), \alpha'_0(t'_j))$, we would otherwise have

$$d_U(\alpha'_0(0), \alpha'_0(t_j)) - d_U(\alpha'_0(0), \alpha'_0(t'_j)) \geq d_U(x'_j, y'_j) - 5\theta_0 \geq 2KE + \theta_0.$$

However, $d_U(\alpha'_0(t_j), \alpha'_0(t'_j)) \leq 2KE$ because of K -discreteness and the projection map to CU being E -coarsely Lipschitz. Hence, the inequality above implies

$$d_U(\alpha'_0(0), \alpha'_0(t_j)) > d_U(\alpha'_0(0), \alpha'_0(t'_j)) + 2KE \geq d_U(\alpha'_0(t_j), \alpha'_0(t'_j)),$$

which contradicts the maximality of t'_j . (Notice that $t'_j \neq k$, and hence $t'_j + 1 \in \{0, \dots, k\}$ because $d_U(\alpha'_0(0), \alpha'_0(t'_j)) + \theta_0 < d_U(\alpha'_0(0), \alpha'_0(t_j)) \leq d_U(\alpha'_0(0), \alpha'_0(k)) + \theta_0$.)

In particular, we get $\text{diam}_{CU}(\beta_j) \leq 2KE + 12\theta_0$, and it is then easy to check that α_0 is $\max\{5KE, 2KE + 12\theta_0\}$ -monotone in U . By replacing K with $\max\{5KE, 2KE + 12\theta_0\}$, we thus have a K -discrete path $\alpha_0 \subset H_\theta(x, y)$ that joins x, y and is K -monotone in U .

We now proceed to the inductive step. Specifically, we fix $\ell \geq 0$ and we assume there exist θ_{ind}, K such that there is a path α in $H_{\theta_{ind}}(x, y)$ connecting x to y such that α is (K, K) -good for $\mathfrak{S}_U^{\ell-1}$.

The coloring: For short, we will say that $V \in \mathfrak{S}$ is A -relevant if $d_U(x, y) \geq A$, see Definition 2.7. Notice that to prove that a path in $H_\theta(x, y)$ is monotone, it suffices to restrict our attention to only those $W \in \mathfrak{S}$ which are, say, $10KE$ -relevant.

By Lemma 2.14, there exists $\chi \geq 0$, bounded by the complexity of \mathcal{X} , and a χ -coloring c of the $10KE$ -relevant elements of \mathfrak{T}_U^ℓ such that $c(V) = c(V')$ only if $V \triangleleft V'$. In other words, the set of $10KE$ -relevant elements of \mathfrak{T}_U^ℓ has the form $\bigsqcup_{i=0}^{\chi-1} c^{-1}(i)$, where $c^{-1}(i)$ is a set of pairwise-transverse relevant elements of \mathfrak{T}_U^ℓ .

Induction hypothesis: Given $p < \chi - 1$, assume by induction (on ℓ and p) that there exist $\theta_p \geq \theta_{ind}, K_p \geq K$, independent of x, y, U , and a path $\alpha_p: \{0, \dots, k\} \rightarrow H_{\theta_p}(x, y)$, joining x, y , that is (K_p, K_p) -good for $\bigsqcup_{i=0}^p c^{-1}(i)$ and good for $\mathfrak{S}_U^{\ell-1}$.

Resolving backtracks in the next color: Let θ_{p+1} be provided by Lemma 4.17 with input θ_p . We will modify α_p to construct a K_{p+1} -discrete path α_{p+1} in $H_{\theta_{p+1}}(x, y)$, for some $K_{p+1} \geq K_p$, that joins x, y and is (K_{p+1}, K_{p+1}) -good in $\bigsqcup_{i=0}^{p+1} c^{-1}(i) \cup \mathfrak{S}_U^{\ell-1}$.

Notice that we can restrict our attention to the set \mathcal{C}_{p+1} of $100(K_p E + \theta_p)$ -relevant elements of $c^{-1}(p+1)$.

A point $t \in \{0, \dots, k\}$ is a $(p+1)$ -omen if there exists $V \in \mathcal{C}_{p+1}$ and $t' > t$ so that $d_V(\alpha_p(0), \alpha_p(t)) > d_V(\alpha_p(0), \alpha_p(t')) + 5K_p E$. If α_p has no $(p+1)$ -omens, then we can take $\alpha_{p+1} = \alpha_p$, since α_p is good in each V with $c(V) < p+1$. Therefore, suppose that there is a $(p+1)$ -omen, let t_0 be the minimal $(p+1)$ -omen, witnessed by $V_0 \in \mathcal{C}_{p+1}$. We can assume that t_0 satisfies $d_{V_0}(\{x, y\}, \alpha_p(t_0)) > 10K_p E$. Let $t'_0 > t_0$ be maximal so that $d_{V_0}(\alpha_p(0), \alpha_p(t_0)) > d_{V_0}(\alpha_p(0), \alpha_p(t'_0))$. In particular $d_{V_0}(y, \alpha_p(t'_0)) \geq 10E$.

Let $x'_0 = \alpha_0(t_0)$ and $y'_0 = \alpha_0(t'_0)$. Inductively, define t_j as the minimal $(p+1)$ -omen, witnessed by $V_j \in \mathcal{C}_{p+1}$, with $t_j \geq t'_{j-1}$, if such t_j exists and let t'_j be maximal so that $d_{V_j}(\alpha_p(0), \alpha_p(t_j)) > d_{V_j}(\alpha_p(0), \alpha_p(t'_j))$ and $d_{V_j}(y, \alpha_p(t'_j)) > 10E$. We can assume that t_j satisfies $d_{V_j}(\{x, y\}, \alpha_p(t_j)) > 10K_p E$. Also, let $x'_j = \alpha_p(t_j), y'_j = \alpha_p(t'_j)$.

Let β_j be a path in $H_{\theta_p}(x'_j, y'_j)$ joining x'_j to y'_j that is (K_p, K_p) -good for each relevant V with $c(V) \leq p$ and each relevant $V \in \mathfrak{S}_U^{\ell-1}$. Such paths can be constructed by induction. By Lemma 4.17 β_j lies in $H_{\theta_{p+1}}(x, y)$. Let α_{p+1} be obtained from α_p by replacing each $\alpha_p(\{t_j, \dots, t'_j\})$ with β_j . Clearly, α_{p+1} connects x to y , is K_p -discrete, and is contained in $H_{\theta_{p+1}}(x, y)$.

We observe that the same argument as in the case $\ell = 0$ gives $d_{V_j}(x'_j, y'_j) \leq 2K_p E + 10\theta_p$.

Verification that α_{p+1} is good for current colors: We next check that each β_j is $10^3(K_p E + \theta_p)$ -monotone in each $W \in \bigsqcup_{i=0}^{p+1} c^{-1}(i)$. We have to consider the following cases. (We can and shall assume below W is $100(K_p E + \theta_p)$ -relevant.)

- If $W \sqsubseteq V_j$, then $W = V_j$, since $\ell_W = \ell_{V_j}$. Since the projections on CW of the endpoints of the straight path β_j coarsely coincide, β_j is $(2K_p E + 12\theta_p)$ -monotone in W . (See the case $\ell = 0$.)
- Suppose $V_j \sqsubset W$. We claim that the projections of the endpoints of β_j lie at a uniformly bounded distance in CW .

We claim that $\rho_W^{V_j}$ has to be E -close to either $[x, x'_j]_W$ or $[y'_j, y]_W$. In fact, if this was not the case, we would have

$$d_{V_j}(x, y) \leq d_{V_j}(x, x'_j) + d_{V_j}(x'_j, y'_j) + d_{V_j}(y'_j, y) \leq 2E + 2K_p E + 10\theta_p,$$


where we applied bounded geodesic image (Definition 1.1.(7)) to the first and last terms.

This is a contradiction with V_j being $100(K_p E + \theta_p)$ -relevant.

Suppose for a contradiction that $d_W(x'_j, y'_j) \geq 500(K_p E + \theta_p)$. Suppose first that $\rho_W^{V_j}$ is E -close to $[x, x'_j]_W$. Then, by monotonicity, $\rho_W^{V_j}$ is E -far from $[\alpha_p(t'_j), y]_W$. By bounded geodesic image, this contradicts $d_{V_j}(y, \alpha_p(t'_j)) \geq 10E$. If instead $\rho_W^{V_j}$ is E -close to $[y'_j, y]_W$, then by bounded geodesic image we have $d_{V_j}(x, \alpha_p(t_j)) \leq E$, contradicting that t_j is an omen witnessed by V_j . See Figure 3.

Hence $d_W(x'_j, y'_j) \leq 500(K_p E + \theta_p)$ and β_j is $10^3(K_p E + \theta_p)$ -monotone in W .

- Suppose $W \not\sqsubset V_j$. We again claim that the projections of the endpoints of β_j are uniformly close in CW , by showing that they both coarsely coincide with $\rho_W^{V_j}$. Since


 FIGURE 3. The situation in CW .

V_j is relevant, either $d_{V_j}(x, \rho_{V_j}^W) \geq E$ or $d_{V_j}(y, \rho_{V_j}^W) \geq E$. Thus, by consistency, $d_W(\rho_W^{V_j}, \{x, y\}) \leq E$. Suppose for a contradiction, that $d_W(x'_j, y'_j) > 100(K_p E + \theta_p)$. We consider separately the cases where $d_W(x, \rho_W^{V_j}) \leq E$ and $d_W(y, \rho_W^{V_j}) \leq E$

First, suppose that $d_W(x, \rho_W^{V_j}) \leq E$. Then $d_W(y, \rho_W^{V_j}) \geq 10K_p E - E > E$, so by consistency, $d_{V_j}(y, \rho_{V_j}^W) \leq E$. If $d_{V_j}(x, \{x'_j, y'_j\}) > E$, then consistency implies that $d_W(x'_j, \rho_W^{V_j}) \leq E$ and $d_W(y'_j, \rho_W^{V_j}) \leq E$, whence $d_W(x'_j, y'_j) \leq 2E$, a contradiction. If $d_{V_j}(x, \{x'_j, y'_j\}) \leq E$, then since $d_{V_j}(x'_j, y'_j) \leq 2K_p E + 10\theta_p$, we have $d_{V_j}(x, x'_j) \leq 5K_p E + 10\theta_p$, contradicting that, since t_j was a $(p+1)$ -omen witnessed by V_j , we must have $d_{V_j}(x, x'_j) > 5K_p E$.

Second, suppose $d_W(y, \rho_W^{V_j}) \leq E$. Then by relevance of W and consistency, $d_{V_j}(x, \rho_{V_j}^W) \leq E$. As above, we have $d_{V_j}(x'_j, x) > 5K_p E + 10\theta_p$, so $d_{V_j}(x, \{x'_j, y'_j\}) > K_p E > 3E$ (since $d_{V_j}(x'_j, y'_j) \leq 2K_p E + 10\theta_p$ and we may assume $K_p > 3$), so $d_{V_j}(\rho_{V_j}^W, \{x'_j, y'_j\}) > E$. Thus, by consistency, $\pi_W(x'_j), \pi_W(y'_j)$ both lie at distance at most E from $\rho_W^{V_j}$, whence $d_W(x'_j, y'_j) \leq 3E$.

- Finally, suppose that $W \perp V_j$. Then either $c(W) < c(V_j)$ and β_j is K_p -monotone in W , or W is irrelevant.

Hence, each β_j is $10^3(K_p E + \theta_p)$ -monotone in each $W \in c^{-1}(\{0, \dots, p+1\})$. Moreover, our above choice of β_j ensures that β_j is K_p -monotone in each $V \in \mathfrak{S}_U^{\ell-1}$.

Verification that α_{p+1} is monotone: Suppose that there exist t, t' such that $t < t'$ and $d_V(\alpha_{p+1}(0), \alpha_{p+1}(t)) > d_V(\alpha_{p+1}(0), \alpha_{p+1}(t')) + 10^4(K_p E + \theta_p)$ for some $V \in c^{-1}(\{0, \dots, p+1\}) \cup \mathfrak{S}_U^{\ell-1}$. We can assume $t, t' \notin \cup_i (t_i, t'_i)$. Indeed, if $t \in (t_i, t'_i)$ (respectively, $t' \in (t_j, t'_j)$), then since all β_m are $10^3(K_p E + \theta_p)$ -monotone, we can replace t with t'_i (respectively, t' with t_j). After such a replacement, we still have $d_V(\alpha_{p+1}(0), \alpha_{p+1}(t)) > d_V(\alpha_{p+1}(0), \alpha_{p+1}(t')) + 5K_p E$.

Let i be maximal so that $t'_i \leq t$ (or let $i = -1$ if no such t'_i exists). By definition of t_{i+1} , we have $t_{i+1} \leq t$, and hence $t_{i+1} = t$. But then $t'_{i+1} > t'$, which is not the case.

Conclusion: Continuing the above procedure while $p < \chi$ produces the desired path α_χ which is (K_χ, K_χ) -good for \mathfrak{S}_U^ℓ . In particular, when $U = S$ is \sqsubseteq -maximal and ℓ is the length of a maximal \sqsubseteq -chain, the proposition follows. \square

4.4. Upper and lower distance bounds. We now state and prove the remaining lemmas needed to complete the proof of Theorem 4.5 and 4.4.

Lemma 4.18 (Upper bound). *For every K, s there exists r with the following property. Let $\alpha: \{0, \dots, n\} \rightarrow \mathcal{X}$ be a K -monotone, (r, K) -proper discrete path connecting x to y . Then*

$$|\alpha| - 1 \leq \sum_{W \in \mathfrak{S}} \{\{d_W(x, y)\}\}_s.$$

Proof. Let $r = r(K, E, s)$ be large enough that, for any $a, b \in \mathcal{X}$, if $d_{\mathcal{X}}(a, b) \geq r$, then there exists $W \in \mathfrak{S}$ so that $d_W(a, b) \geq 100KEs$. This r is provided by Definition 1.1.(9).

For $0 \leq j \leq n-1$, choose $V_j \in \mathfrak{S}$ so that $d_{V_j}(\alpha(j), \alpha(j+1)) \geq 100KEs$. By monotonicity of α in V_j , for any $j' > j$ we have

$$d_{V_j}(\alpha(0), \alpha(j')) \geq d_{V_j}(\alpha(0), \alpha(j)) + 50KEs.$$

It follows by induction on $j \leq n$ that $\sum_{W \in \mathfrak{S}} \{\{d_W(\alpha(0), \alpha(j))\}\}_s \geq \min\{j, n-1\}$. \square

Lemma 4.19 (Lower bound). *There exists s_0 such that for all $s \geq s_0$, there exists C with the following property.*

$$d_{\mathcal{X}}(x, y) \geq \frac{1}{C} \sum_{W \in \mathfrak{S}} \{\{d_W(x, y)\}\}_s.$$

Proof. From Corollary 4.16, we obtain a K -discrete path $\alpha: \{0, n\} \rightarrow \mathcal{X}$ joining x, y and having the property that the (coarse) path $\pi_V \circ \alpha: \{0, \dots, n\} \rightarrow \mathcal{C}V$ lies in the K -neighborhood of a geodesic from $\pi_V(x)$ to $\pi_V(y)$. Moreover, α is K -efficient, by the same corollary.

Fix $s_0 \geq 10^3KE$. A *checkpoint* for x, y in $V \in \mathfrak{S}$ is a ball Q in $\mathcal{C}V$ so that $\pi_V \circ \alpha$ intersects Q and $d_V(\{x, y\}, Q) \geq 10KE + 1$. Note that any ball of radius $10KE$ centered on a geodesic from $\pi_V(x)$ to $\pi_V(y)$ is a checkpoint for x, y in V , provided it is sufficiently far from $\{x, y\}$.

For each $V \in \mathbf{Rel}(x, y, 10^3KE)$, choose a set \mathfrak{C}_V of $\left\lceil \frac{d_V(x, y)}{10} \right\rceil$ checkpoints for x, y in V , subject to the requirement that $d_V(C_1, C_2) \geq 10KE$ for all distinct $C_1, C_2 \in \mathfrak{C}_V$. For each $V \in \mathbf{Rel}(x, y, 10^3KE)$, we have $10|\mathfrak{C}_V| \geq d_V(x, y)$, so

$$\sum_{V \in \mathfrak{S}} |\mathfrak{C}_V| \geq \frac{1}{10} \sum_{W \in \mathfrak{S}} \{\{d_W(x, y)\}\}_{s_0}.$$

Each $j \in \{0, \dots, n\}$ is a *door* if there exists $V \in \mathbf{Rel}(x, y, 10^3KE)$ and $C \in \mathfrak{C}_V$ such that $\pi_V(\alpha(j)) \in C$ but $\pi_V(\alpha(j-1)) \notin C$. The *multiplicity* of a door j is the cardinality of the set $\mathcal{M}(j)$ of $V \in \mathbf{Rel}(x, y, 10^3KE)$ for which there exists $C \in \mathfrak{C}_V$ with $\pi_V(\alpha(j)) \in C$ and $\pi_V(\alpha(j-1)) \notin C$. Since \mathfrak{C}_V is a set of pairwise-disjoint checkpoints, j is a door for at most one element of \mathfrak{C}_V , for each V . Hence the multiplicity of j is precisely the total number of checkpoints in $\cup_{V \in \mathbf{Rel}(x, y, 10^3KE)} \mathfrak{C}_V$ for which j is a door.

We claim that the set $\mathcal{M}(j)$ does not contain a pair of transverse elements. Indeed, suppose that $U, V \in \mathcal{M}(j)$, satisfy $U \pitchfork V$. Let $Q_V \in \mathfrak{C}_V, Q_U \in \mathfrak{C}_U$ be the checkpoints containing $\pi_V(\alpha(j)), \pi_U(\alpha(j))$ respectively, so that $d_U(\alpha(j), \{x, y\}), d_V(\alpha(j), \{x, y\}) \geq 10KE + 1 > 10E$, contradicting Corollary 2.4. Lemma 2.2 thus gives $|\mathcal{M}(j)| \leq \chi$. Now, $|\alpha|$ is at least the number of doors in $\{0, \dots, n\}$, whence $|\alpha| \geq \frac{1}{\chi} \sum_{V \in \mathfrak{S}} |\mathfrak{C}_V|$. Since α is K -efficient, we obtain

$$d_{\mathcal{X}}(x, y) \geq \frac{1}{10\chi K} \sum_{W \in \mathfrak{S}} \{\{d_W(x, y)\}\}_{s_0}.$$

For $s \geq s_0$, $\sum_{W \in \mathfrak{S}} \{\{d_W(x, y)\}\}_s \leq \sum_{W \in \mathfrak{S}} \{\{d_W(x, y)\}\}_{s_0}$, so the claim follows. \square

5. HIERARCHICAL QUASICONVEXITY AND GATES

We now introduce the notion of hierarchical quasiconvexity, which is essential for the discussion of product regions, the combination theorem of Section 8, and in the forthcoming [DHS17].

Definition 5.1 (Hierarchical quasiconvexity). 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 3.1) satisfies $d(x, \mathcal{Y}) \leq k(\kappa)$.

Remark 5.2. Note that condition (2) in the above definition is equivalent to: For every $\kappa > 0$ and every point $x \in \mathcal{X}$ satisfying $d_U(\pi_U(x), \pi_U(\mathcal{Y})) \leq \kappa$ for all $U \in \mathfrak{S}$, has the property that $d(x, \mathcal{Y}) \leq k(\kappa)$.

Lemma 5.3. For each Q there exists κ so that the following holds. Let $\mathcal{Y} \subseteq \mathcal{X}$ be such that $\pi_V(\mathcal{Y})$ is Q -quasiconvex for each $V \in \mathfrak{S}$. Let $x \in \mathcal{X}$ and, for each $V \in \mathfrak{S}$, let $p_V \in \pi_V(\mathcal{Y})$ satisfy $d_V(x, p_V) \leq d_V(x, \mathcal{Y}) + 1$. Then (p_V) is κ -consistent.

Proof. For each V , choose $y_V \in \mathcal{Y}$ so that $\pi_V(y_V) = p_V$.

Suppose that $V \triangleleft W$ or $V \sqsubseteq W$. By Lemma 2.6 and Theorem 3.1, there exists $z \in \mathcal{X}$ so that for all $U \in \mathfrak{S}$, the projection $\pi_U(z)$ lies C -close to each of the geodesics $[\pi_U(x), \pi_U(y_V)]$, $[\pi_U(x), \pi_U(y_W)]$, and $[\pi_U(y_W), \pi_U(y_V)]$, where C depends on \mathcal{X} . Hence $d_V(p_V, z)$ and $d_W(p_W, z)$ are uniformly bounded, by quasiconvexity of $\pi_V(\mathcal{Y})$ and $\pi_W(\mathcal{Y})$.

Suppose that $V \triangleleft W$. Since the tuple $(\pi_U(z))$ is consistent, either y_V lies uniformly close in $\mathcal{C}V$ to ρ_V^W , or the same holds with V and W interchanged, as required. Suppose that $V \sqsubseteq W$. Suppose that $d_W(p_W, \rho_W^V)$ is sufficiently large, so that we have to bound $\text{diam}_V(\rho_V^W(p_W) \cup p_V)$. Since $d_W(z, p_W)$ is uniformly bounded, $d_W(z, \rho_W^V)$ is sufficiently large that consistency ensures that $\text{diam}_V(\rho_V^W(\pi_W(z)) \cup \pi_V(z))$ is uniformly bounded. Since any geodesic from p_W to z lies far from ρ_W^V , the sets $\rho_V^W(\pi_W(z))$ and $\rho_V^W(p_V)$ coarsely coincide. Since $\pi_V(z)$ coarsely coincides with p_V by construction of z , we have the required bound. Hence the tuple with V -coordinate p_V is κ -consistent for uniform κ . \square

Definition 5.4 (Gate). A coarsely Lipschitz map $\mathfrak{g}_\mathcal{Y}: \mathcal{X} \rightarrow \mathcal{Y}$ is called a *gate map* if for each $x \in \mathcal{X}$ it satisfies: $\mathfrak{g}_\mathcal{Y}(x)$ is a point $y \in \mathcal{Y}$ such that for all $V \in \mathfrak{S}$, the set $\pi_V(y)$ (uniformly) coarsely coincides with the projection of $\pi_V(x)$ to the $k(0)$ -quasiconvex set $\pi_V(\mathcal{Y})$. The point $\mathfrak{g}(x)$ is called the *gate of x in \mathcal{Y}* . The uniqueness axiom implies that when such a map exists it is coarsely well-defined.

We first establish that, as should be the case for a (quasi)convexity property, one can coarsely project to hierarchically quasiconvex subspaces. The next lemma shows that gates exist for k -hierarchically quasiconvex subsets.

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

Proof. For each $V \in \mathfrak{S}$, let $p_V \in \pi_V(\mathcal{Y})$ satisfy $d_V(x, p_V) \leq d_V(x, \mathcal{Y}) + 1$. Then (p_V) is κ -consistent for some κ independent of x by Lemma 5.3. (Note that (p_V) is admissible by construction.)

Theorem 3.1 and the definition of hierarchical quasiconvexity combine to supply $y' \in \mathcal{N}_{k(\kappa)}(\mathcal{Y})$ with the desired projections to all $V \in \mathfrak{S}$; this point lies at distance $k(\kappa)$ from some $y \in \mathcal{Y}$ with the desired property.

We now check that this map is coarsely Lipschitz. Let $x_0, x_n \in \mathcal{X}$ be joined by a uniform quasigeodesic γ . By sampling γ , we obtain a discrete path $\gamma': \{0, \dots, n\} \rightarrow \mathcal{X}$ such that $d_{\mathcal{X}}(\gamma'(i), \gamma'(i+1)) \leq K$ for $0 \leq i \leq n-1$, where K depends only on \mathcal{X} , and such that $\gamma'(0) = x_0, \gamma'(n) = x_n$. Observe that $d_{\mathcal{X}}(\mathfrak{g}_{\mathcal{Y}}(x_0), \mathfrak{g}_{\mathcal{Y}}(x_n)) \leq \sum_{i=0}^{n-1} d_{\mathcal{X}}(\mathfrak{g}_{\mathcal{Y}}(\gamma'(i)), \mathfrak{g}_{\mathcal{Y}}(\gamma'(i+1)))$, so it suffices to exhibit C such that $d_{\mathcal{X}}(\mathfrak{g}_{\mathcal{Y}}(x), \mathfrak{g}_{\mathcal{Y}}(x')) \leq C$ whenever $d_{\mathcal{X}}(x, x') \leq K$. But if $d_{\mathcal{X}}(x, x') \leq K$, then each $d_U(x, x') \leq K'$ for some uniform K' , by Definition 1.1.(1), whence the claim follows from the fact that each $\mathcal{C}U \rightarrow \pi_U(\mathcal{Y})$ is coarsely Lipschitz (with constant depending only on δ and $k(0)$) along with the uniqueness axiom (Definition 1.1.(9)). \square

5.1. Hierarchically quasiconvex subspaces are hierarchically hyperbolic.

Proposition 5.6. *Let $\mathcal{Y} \subseteq \mathcal{X}$ be a hierarchically k -quasiconvex subset of the hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$. Then $(\mathcal{Y}, \mathfrak{d})$ is a hierarchically hyperbolic space, where \mathfrak{d} is the metric inherited from \mathcal{X} .*

Proof. There exists K so that any two points in \mathcal{Y} are joined by a uniform quasigeodesic. Indeed, any two points in \mathcal{Y} are joined by a hierarchy path in \mathcal{X} , which must lie uniformly close to \mathcal{Y} .

We now define a hierarchically hyperbolic structure. For each U , let $r_U: \mathcal{C}U \rightarrow \pi_U(\mathcal{Y})$ be the coarse projection, which exists by quasiconvexity. The index set is \mathfrak{S} , and the associated hyperbolic spaces are the various $\mathcal{C}U$. For each U , define a projection $\pi'_U: \mathcal{C}U \rightarrow \mathcal{C}U$ by $\pi'_U = r_U \circ \pi_U$, and for each non-orthogonal pair $U, V \in \mathfrak{S}$, the corresponding relative projection $\mathcal{C}U \rightarrow \mathcal{C}V$ is given by $r_V \circ \rho_V^U$. All of the requirements of Definition 1.1 involving only the various relations on \mathfrak{S} are obviously satisfied, since we have only modified the projections. The consistency inequalities continue to hold since each r_U is uniformly coarsely Lipschitz. The same is true for bounded geodesic image and the large link lemma. Partial realization holds by applying the map $\mathfrak{g}_{\mathcal{Y}}$ to points constructed using partial realization in $(\mathcal{X}, \mathfrak{S})$. \square

Remark 5.7 (Alternative hierarchically hyperbolic structures). In the above proof, one can replace each $\mathcal{C}U$ by a thickening $\mathcal{C}U_{\mathcal{Y}}$ of $\pi_U(\mathcal{Y})$ (this set is quasiconvex; the thickening is to make a hyperbolic geodesic space). This yields a hierarchically hyperbolic structure with coarsely surjective projections.

5.2. Standard product regions. In this section, we describe a class of hierarchically quasiconvex subspaces called *standard product regions* that will be useful in future applications. We first recall a construction from [BHS17a, Section 13].

Definition 5.8 (Nested partial tuple). Recall $\mathfrak{S}_U = \{V \in \mathfrak{S} : V \sqsubseteq U\}$. Fix $\kappa \geq \kappa_0$ and let \mathbf{F}_U be the set of κ -consistent tuples in $\prod_{V \in \mathfrak{S}_U} 2^{\mathcal{C}V}$.

Definition 5.9 (Orthogonal partial tuple). Let $\mathfrak{S}_U^{\perp} = \{V \in \mathfrak{S} : V \perp U\} \cup \{A\}$, where A is a \sqsubseteq -minimal element A such that $V \sqsubseteq A$ for all $V \perp U$. Fix $\kappa \geq \kappa_0$, let \mathbf{E}_U be the set of κ -consistent tuples in $\prod_{V \in \mathfrak{S}_U^{\perp} - \{A\}} 2^{\mathcal{C}V}$.

Construction 5.10 (Product regions in \mathcal{X}). Given \mathcal{X} and $U \in \mathfrak{S}$, there are coarsely well-defined maps $\phi^{\sqsubseteq}, \phi^{\perp}: \mathbf{F}_U, \mathbf{E}_U \rightarrow \mathcal{X}$ whose images are hierarchically quasiconvex and which extend to a coarsely well-defined map $\phi_U: \mathbf{F}_U \times \mathbf{E}_U \rightarrow \mathcal{X}$ with hierarchically quasiconvex image. Indeed, for each $(\vec{a}, \vec{b}) \in \mathbf{F}_U \times \mathbf{E}_U$, and each $V \in \mathfrak{S}$, define the co-ordinate $(\phi_U(\vec{a}, \vec{b}))_V$ 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 $V \not\sqsubseteq U$, then $(\phi_U(\vec{a}, \vec{b}))_V = \rho_V^U$. Finally, if $U \sqsubseteq V$, and $U \neq V$, let $(\phi_U(\vec{a}, \vec{b}))_V = \rho_V^U$.

We now verify that the tuple $\phi_U(\vec{a}, \vec{b})$ is consistent. If $W, V \in \mathfrak{S}$, and either V or W is transverse to U , then the consistency inequality involving W and V is satisfied in view of Proposition 1.8. The same holds if $U \sqsubseteq W$ or $U \sqsubseteq V$. Hence, it remains to consider the cases where V and W are each either nested into or orthogonal to U : if $V, W \sqsubseteq U$ or $V, W \perp U$ then consistency holds by assumption; otherwise, up to reversing the roles of V and W we have $V \sqsubseteq U$ and $W \perp U$, in which case $V \perp W$ and there is nothing to check. Theorem 3.1 thus supplies the map $\phi_U: \mathbf{F}_U \times \mathbf{E}_U \rightarrow \mathcal{X}$. Fixing any $e \in \mathbf{E}_U$ yields a map $\phi_U^\sqsubseteq: \mathbf{F}_U \times \{e\} \rightarrow \mathcal{X}$, and ϕ^\perp is defined analogously. Note that these maps depend on choices of basepoints in $\mathbf{E}_U, \mathbf{F}_U$.

Where it will not introduce confusion (e.g., where the basepoints are understood or immaterial), we abuse notation and regard $\mathbf{F}_U, \mathbf{E}_U$ as subspaces of \mathcal{X} , i.e., $\mathbf{F}_U = \text{im } \phi_U^\sqsubseteq, \mathbf{E}_U = \text{im } \phi^\perp$.

Proposition 5.11. *When $\mathbf{E}_U, \mathbf{F}_U \subset \mathcal{X}$ are endowed with the subspace metric \mathbf{d} , the spaces $(\mathbf{F}_U, \mathfrak{S}_U)$ and $(\mathbf{E}_U, \mathfrak{S}_U^\perp)$ are hierarchically hyperbolic; if U is not \sqsubseteq -maximal, then their complexity is strictly less than that of $(\mathcal{X}, \mathfrak{S})$. Moreover, ϕ_U^\sqsubseteq and ϕ^\perp determine hieromorphisms $(\mathbf{F}_U, \mathfrak{S}_U), (\mathbf{E}_U, \mathfrak{S}_U^\perp) \rightarrow (\mathcal{X}, \mathfrak{S})$.*

Proof. For each $V \sqsubseteq U$ or $V \perp U$, the associated hyperbolic space $\mathcal{C}V$ is exactly the one used in the hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{S})$. For A , use an appropriate thickening \mathcal{C}^*A of $\pi_A(\text{im } \phi^\perp)$ to a hyperbolic geodesic space. All of the projections $\mathbf{F}_U \rightarrow \mathcal{C}V, V \in \mathfrak{S}_U$ and $\mathbf{E}_U \rightarrow \mathcal{C}V, V \in \mathfrak{S}_U^\perp$ are as in $(\mathcal{X}, \mathfrak{S})$ (for A , compose with a quasi-isometry $\pi_A(\text{im } \phi^\perp) \rightarrow \mathcal{C}^*A$). Observe that $(\mathbf{F}_U, \mathfrak{S}_U)$ and $(\mathbf{E}_U, \mathfrak{S}_U^\perp)$ are hierarchically hyperbolic (this can be seen using a simple version of the proof of Proposition 5.6). If U is not \sqsubseteq -maximal in \mathfrak{S} , then neither is A , whence the claim about complexity.

The hieromorphisms are defined by the inclusions $\mathfrak{S}_U, \mathfrak{S}_U^\perp \rightarrow \mathfrak{S}$ and, for each $V \in \mathfrak{S}_U \cup \mathfrak{S}_U^\perp$, the identity $\mathcal{C}V \rightarrow \mathcal{C}V$, unless $V = A$, in which case we use $\mathcal{C}^*A \rightarrow \pi_A(\text{im } \phi^\perp) \hookrightarrow \mathcal{C}A$. These give hieromorphisms by definition. \square

Remark 5.12 (Dependence on A). Note that A need not be the unique \sqsubseteq -minimal element of \mathfrak{S} into which each $V \perp U$ is nested; the axioms don't require uniqueness of such \sqsubseteq -minimal elements. Observe that \mathbf{E}_U (as a set and as a subspace of \mathcal{X}) is defined independently of the choice of A . It is the hierarchically hyperbolic structure from Proposition 5.11 that a priori depends on A . However, note that $A \not\sqsubseteq U$, since there exists $V \sqsubseteq A$ with $V \perp U$, and we cannot have $V \sqsubseteq U$ and $V \perp U$ simultaneously. Likewise, $A \not\perp U$ by definition. Finally, if $U \sqsubseteq A$, then the axioms guarantee the existence of B , properly nested into A , into which each $V \perp U$ is nested, contradicting \sqsubseteq -minimality of A . Hence $U \not\sqsubseteq A$. It follows that $\pi_A(\mathbf{E}_U)$ is bounded — it coarsely coincides with ρ_A^U . Thus the hierarchically hyperbolic structure on \mathbf{E}_U , and the hieromorphism structure of ϕ^\perp , is actually essentially canonical: we can take the hyperbolic space associated to the \sqsubseteq -maximal element to be a point, whose image in each of the possible choices of A must coarsely coincide with ρ_A^U .

Remark 5.13 (Orthogonality and product regions). If $U \perp V$, then we have $\mathbf{F}_U \subseteq \mathbf{E}_V$ and $\mathbf{F}_V \subseteq \mathbf{E}_U$, so there is a hierarchically quasiconvex map $\phi_U \times \phi_V: \mathbf{F}_U \times \mathbf{F}_V \rightarrow \mathcal{X}$ extending to $\phi_U \times \phi_U^\perp$ and $\phi_V^\perp \times \phi_V$.

Remark 5.14. Since $\mathbf{F}_U, \mathbf{E}_U$ are hierarchically quasiconvex spaces, Definition 5.4 provides coarse gates $\mathbf{g}_{\mathbf{F}_U}: \mathcal{X} \rightarrow \mathbf{F}_U$ and $\mathbf{g}_{\mathbf{E}_U}: \mathcal{X} \rightarrow \mathbf{E}_U$. These are coarsely the same as the following maps: given $x \in \mathcal{X}$, let \vec{x} be the tuple defined by $x_W = \pi_W(x)$ when $W \sqsubseteq U$ and $x_W = \pi_W(x)$ when $W \perp U$ and ρ_W^U otherwise. Then \vec{x} is consistent and coarsely equals $\mathbf{g}_{\mathbf{F}_U \times \mathbf{E}_U}(x)$.

Definition 5.15 (Standard product region). For each $U \in \mathfrak{S}$, let $P_U = \text{im } \phi_U$, which is coarsely $\mathbf{F}_U \times \mathbf{E}_U$. We call this the *standard product region* in \mathcal{X} associated to U .

The next proposition follows from the definition of the product regions and the fact that, if $U \sqsubseteq V$, then ρ_W^U, ρ_W^V coarsely coincide whenever $V \sqsubseteq W$ or $V \pitchfork W$ and $U \not\sqsubseteq W$, which holds by Definition 1.1.(4).

Proposition 5.16 (Parallel copies). *There exists $\nu \geq 0$ such that for all $U \in \mathfrak{S}$, all $V \in \mathfrak{S}_U$, and all $u \in \mathbf{E}_U$, there exists $v \in \mathbf{E}_V$ so that $\phi_V(\mathbf{F}_V \times \{v\}) \subseteq \mathcal{N}_\nu(\phi_U(\mathbf{F}_U \times \{u\}))$.*

5.2.1. *Hierarchy paths and product regions.* Recall that a D -hierarchy path γ in \mathcal{X} is a (D, D) -quasigeodesic $\gamma : I \rightarrow \mathcal{X}$ such that $\pi_U \circ \gamma$ is an unparameterized (D, D) -quasigeodesic for each $U \in \mathfrak{S}$, and that Theorem 4.4 provides $D \geq 1$ so that any two points in \mathcal{X} are joined by a D -hierarchy path. In this section, we describe how hierarchy paths interact with standard product regions.

In the next proposition and lemma, given $x, y \in \mathcal{X}$, we declare $V \in \mathfrak{S}$ to be *relevant* (for x, y) if $d_V(x, y) \geq 200DE$.

Proposition 5.17 (“Active” subpaths). *There exists $\nu \geq 0$ so that for all $x, y \in \mathcal{X}$, all $V \in \mathfrak{S}$ with V relevant for (x, y) , and all D -hierarchy paths γ joining x to y , there is a subpath α of γ with the following properties:*

- (1) $\alpha \subset \mathcal{N}_\nu(P_V)$;
- (2) $\pi_U|_\gamma$ is coarsely constant on $\gamma - \alpha$ for all $U \in \mathfrak{S}_V \cup \mathfrak{S}_V^\perp$.

Proof. We may assume $\gamma : \{0, n\} \rightarrow \mathcal{X}$ is a $2D$ -discrete path. Let $x_i = \gamma(i)$ for $0 \leq i \leq n$. Let $S \in \mathfrak{S}$ be the \sqsubseteq -maximal element. Since the proposition holds trivially for $V = S$, assume $V \sqsubset S$.

First consider the case where V is \sqsubseteq -maximal among relevant elements of \mathfrak{S} . Lemma 5.18 provides $\nu'' \geq 0$, independent of x, y , and also provides $i \leq n$, such that $d_S(x_i, \rho_S^V) \leq \nu''$. Let i be minimal with this property and let i' be maximal with this property. Observe that there exists $\nu' \geq \nu''$, depending only on ν'' and the (uniform) monotonicity of γ in \mathcal{CS} , such that $d_S(x_j, \rho_S^V) \leq \nu'$ for $i \leq j \leq i'$.

For $j \in \{i, \dots, i'\}$, let $x'_j = \mathbf{g}_{P_V}(x_j)$. Let $U \in \mathfrak{S}$. By definition, if $U \sqsubseteq V$ or $U \perp V$, then $\pi_U(x_j)$ coarsely coincides with $\pi_U(x'_j)$, while $\pi_U(x'_j)$ coarsely coincides with ρ_U^V if $V \sqsubseteq U$ or $V \pitchfork U$. We claim that there exist i_1, i'_1 with $i \leq i_1 \leq i'_1 \leq i'$ such that for $i_1 \leq j \leq i'_1$ and $U \in \mathfrak{S}$ with $V \sqsubseteq U$ or $U \pitchfork V$, the points $\pi_U(x_j)$ and $\pi_U(x'_j)$ coarsely coincide; this amounts to claiming that $\pi_U(x_j)$ coarsely coincides with ρ_U^V .

If $V \sqsubseteq U$ and some geodesic σ in \mathcal{CU} from $\pi_U(x)$ to $\pi_U(y)$ fails to pass through the E -neighborhood of ρ_U^V , then bounded geodesic image shows that $\rho_V^U(\sigma)$ has diameter at most E . On the other hand, consistency shows that the endpoints of $\rho_V^U(\sigma)$ coarsely coincide with $\pi_V(x)$ and $\pi_V(y)$, contradicting that V is relevant. Thus σ passes through the E -neighborhood of ρ_U^V . Maximality of V implies that U is not relevant, so that $\pi_V(x), \pi_V(y)$, and $\pi_V(x_j)$ all coarsely coincide, whence $\pi_V(x_j)$ coarsely coincides with ρ_V^U .

If $U \pitchfork V$ and U is not relevant, then $\pi_U(x_j)$ coarsely coincides with both $\pi_U(x)$ and $\pi_U(y)$, each of which coarsely coincides with ρ_U^V , for otherwise we would have $d_V(x, y) \leq 2E$ by consistency and the triangle inequality, contradicting that V is relevant. If $U \pitchfork V$ and U is relevant, then, by consistency, we can assume that $\pi_U(y), \rho_U^V$ coarsely coincide, as do $\pi_V(x), \rho_V^U$. Either $\pi_U(x_j)$ coarsely equals ρ_U^V , or $\pi_V(x_j)$ coarsely equals $\pi_V(x)$, again by consistency. If $d_V(x, x_j) \leq 10E$ or $d_V(y, x_j) \leq 10E$, discard x_j . Our discreteness assumption and the fact that V is relevant imply that there exist $i_1 \leq i'_1$ between i and i' so that x_j is not discarded for $i_1 \leq j \leq i'_1$. For such j , the distance formula now implies that $d(x_j, x'_j)$ is bounded by a constant ν independent of x, y .

We thus have i_1, i'_1 such that $x_j \in \mathcal{N}_\nu(P_V)$ for $i \leq j \leq i'$ and $x_j \notin \mathcal{N}_\nu(P_V)$ for $j < i$ or $j > i'$, provided V is \sqsubseteq -maximal relevant. If $W \sqsubseteq V$ and W is relevant, and there

is no relevant $W' \neq W$ with $W \sqsubseteq W' \sqsubseteq V$, then we may apply the above argument to $\gamma' = \mathbf{g}_{P_V}(\gamma|_{i, \dots, i'})$ to produce a subpath of γ' lying ν -close to $P_W \subseteq P_V$, and hence a subpath of γ lying 2ν -close to P_W . Finiteness of the complexity (Definition 1.1.(5)) then yields assertion (1). Assertion (2) is immediate from our choice of i_1, i'_1 . \square

Lemma 5.18. *There exists $\nu' \geq 0$ so that for all $x, y \in \mathcal{X}$, all relevant $V \in \mathfrak{S}$, and all D -hierarchy paths γ joining x to y , there exists $t \in \gamma$ so that $\mathbf{d}_S(t, \rho_S^V) \leq \nu'$.*

Proof. Let σ be a geodesic in \mathcal{CS} joining the endpoints of $\pi_S \circ \gamma$. Since $\mathbf{d}_V(x, y) \geq 200DE$, the consistency and bounded geodesic image axioms (Definition 1.1.(4),(7)) imply that σ enters the E -neighborhood of ρ_S^V in \mathcal{CS} , whence $\pi_S \circ \gamma$ comes uniformly close to ρ_S^V . \square

6. HULLS

In this section we build ‘‘convex hulls’’ in hierarchically hyperbolic spaces. This construction is motivated by, and generalizes, the concept in the mapping class group called Σ -hull, as defined by Behrstock–Kleiner–Minsky–Mosher [BKMM12]. Recall that given a set A of points in a δ -hyperbolic space H , its *convex hull*, denoted $\text{hull}_H(A)$, is the union of geodesics between pairs of points in this set. We will make use of the fact that the convex hull is 2δ -quasiconvex (since, if $p \in [x, y], q \in [x', y']$, then $[p, q] \subseteq \mathcal{N}_{2\delta}([p, x] \cup [x, x'] \cup [x', q]) \subseteq \mathcal{N}_{2\delta}([y, x] \cup [x, x'] \cup [x', y']$)).

The construction of hulls is based on Proposition 6.3, which generalizes Lemma 4.15; indeed, the construction of hulls in this section generalizes the hulls of pairs of points used in Section 4 to prove the distance formula. The second part of Proposition 6.3 (which is not used in Section 4) relies on the distance formula.

Definition 6.1 (Hull of a set). For each set $A \subset \mathcal{X}$ and $\theta \geq 0$, let $H_\theta(A)$ 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 $\text{hull}_{CW}(A)$. Note that $A \subset H_\theta(A)$.

Lemma 6.2. *There exists θ_0 so that for each $\theta \geq \theta_0$ there exists $k : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and each $A \subseteq \mathcal{X}$, we have that $H_\theta(A)$ is k -hierarchically quasiconvex.*

Proof. For any θ and $U \in \mathfrak{S}$, due to δ -hyperbolicity we have that $\pi_U(H_\theta(A))$ is 2δ -quasiconvex, so we only have to check the condition on realization points.

Let A' be the union of all D_0 -hierarchy paths joining points in A , where D_0 is the constant from Theorem 4.4. Then the Hausdorff distance between $\pi_U(A')$ and $\pi_U(A)$ is bounded by $C = C(\delta, D_0)$ for each $U \in \mathfrak{S}$. Also, $\pi_U(A')$ is $Q = Q(\delta, D_0)$ -quasiconvex. Let κ be the constant from Lemma 5.3, and let $\theta_0 = \theta_e(\kappa)$ be as in Theorem 3.1.

Fix any $\theta \geq \theta_0$, and any $\kappa \geq 0$. Let (b_U) be a κ' -consistent tuple with $b_U \subseteq \mathcal{N}_\theta(\text{hull}_{CU}(A))$ for each $U \in \mathfrak{S}$. Let $x \in \mathcal{X}$ project $\theta_e(\kappa')$ -close to each b_U . We have to find $y \in H_\theta(A)$ uniformly close to x . By Lemma 5.3, (p_U) is κ -consistent, where $p_U \in \text{hull}_{CW}(A)$ satisfies $\mathbf{d}_U(x, p_U) \leq \mathbf{d}_U(x, \text{hull}_{CW}(A)) + 1$. It is readily seen from the uniqueness axiom (Definition 1.1.9) that any $y \in \mathcal{X}$ projecting close to each p_U has the required property, and such a y exists by Theorem 3.1. To check admissibility, note that each p_U lies θ -close to $\text{hull}_{CU}(A)$, which in turn lies uniformly close to $\pi_U(\mathcal{X})$ by quasiconvexity of $\pi_U(\mathcal{X})$. \square

We denote the Hausdorff distance in the metric space Y by $\mathbf{d}_{\text{Haus}, Y}(\cdot, \cdot)$. The next proposition directly generalizes [BKMM12, Proposition 5.2] from mapping class groups to general hierarchically hyperbolic spaces.

Proposition 6.3 (Retraction onto hulls). *For each sufficiently large θ there exists $C \geq 1$ so that for each set $A \subset \mathcal{X}$ there is a (K, K) -coarsely Lipschitz map $r : \mathcal{X} \rightarrow H_\theta(A)$ restricting to the identity on $H_\theta(A)$. Moreover, if $A' \subset \mathcal{X}$ lies at finite Hausdorff distance from A , then $\mathbf{d}_{\mathcal{X}}(r_A(x), r_{A'}(x))$ is C -coarsely Lipschitz in $\mathbf{d}_{\text{Haus}, \mathcal{X}}(A, A')$.*

Proof. By Lemma 6.2, for all sufficiently large θ , $H_\theta(A)$ is hierarchically quasiconvex. Thus, by Lemma 5.5 there exists a map $r: \mathcal{X} \rightarrow H_\theta(A)$, which is coarsely Lipschitz and which is the identity on $H_\theta(A)$.

We now prove the moreover clause. By Definition 1.1.(1), for each W the projections π_W are each coarsely Lipschitz and thus $\mathbf{d}_{Haus, \mathcal{C}W}(\pi_W(A), \pi_W(A'))$ is bounded by a coarsely Lipschitz function of $\mathbf{d}_{Haus, \mathcal{X}}(A, A')$. It is then easy to conclude using the distance formula (Theorem 4.5) and the construction of gates (Definition 5.4) used to produce the map r . \square

6.1. Homology of asymptotic cones. In this subsection we make a digression to study homological properties of asymptotic cones of hierarchically hyperbolic spaces. This subsection is not needed for the proof of distance formula, and in fact we will use the distance formula in a proof.

Using Proposition 6.3, the identical proof as used in [BKMM12, Lemma 5.4] for mapping class groups, yields:

Proposition 6.4. *There exists $\theta_0 \geq 0$ depending only on the constants of the hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ such that for all $\theta, \theta' \geq \theta_0$ there exist K, C , and θ'' such that given two sets $A, A' \subset \mathcal{X}$, then:*

- (1) $\text{diam}(H_\theta(A)) \leq K \text{diam}(A) + C$
- (2) If $A' \subset H_\theta(A)$ then $H_\theta(A') \subset H_{\theta''}(A)$.
- (3) $d_{Haus, \mathcal{X}}(H_\theta(A), H_\theta(A')) \leq K d_{Haus, \mathcal{X}}(A, A') + C$.
- (4) $d_{Haus, \mathcal{X}}(H_\theta(A), H_{\theta'}(A)) \leq C$

Remark 6.5. Proposition 6.4 is slightly stronger than the corresponding [BKMM12, Lemma 5.4], in which A, A' are finite sets and the constants depend on their cardinality. The source of the strengthening is just the observation that hulls in δ -hyperbolic spaces are 2δ -quasiconvex regardless of the cardinality of the set (see [BKMM12, Lemma 5.1]).

It is an easy observation that given a sequence \mathbf{A} of sets $A_n \subset \mathcal{X}$ with bounded cardinality, the retractions to the corresponding hulls $H_\theta(A_n)$ converge in any asymptotic cone, \mathcal{X}_ω , to a Lipschitz retraction from that asymptotic cone to the ultralimit of the hulls, $H(\mathbf{A})$. A general argument, see e.g., [BKMM12, Lemma 6.2] implies that the ultralimit of the hulls is then contractible. The proofs in [BKMM12, Section 6] then apply in the present context using the above Proposition, with the only change needed that the reference to the rank theorem for hierarchically hyperbolic spaces as proven in [BHS17a, Theorem J] must replace the application of [BM08]. In particular, this yields the following two results:

Corollary 6.6. *Let \mathcal{X} be a hierarchically hyperbolic space and \mathcal{X}_ω one of its asymptotic cones. Let $X \subset \mathcal{X}_\omega$ be an open subset and suppose that for any sequence, \mathbf{A} , of finite subsets of X we have $H(\mathbf{A}) \subset X$. Then X is acyclic.*

Corollary 6.7. *If (U, V) is an open pair in \mathcal{X}_ω , then $H_k(U, V) = \{0\}$ for all k greater than the complexity of \mathcal{X} .*

6.2. Relatively hierarchically hyperbolic spaces and the distance formula. In this section, we work in the following context:

Definition 6.8 (Relatively hierarchically hyperbolic spaces). The hierarchical space $(\mathcal{X}, \mathfrak{S})$ is *relatively hierarchically hyperbolic* if there exists δ such that for all $U \in \mathfrak{S}$, either U is \sqsubseteq -minimal or $\mathcal{C}U$ is δ -hyperbolic. If U is \sqsubseteq -minimal and $\mathcal{C}U$ is not hyperbolic, then we insist that π_U is E -coarsely surjective.

Remark 6.9. One could, more generally, only insist that each $\pi_U(\mathcal{X})$ is a uniformly coarsely Lipschitz coarse retract. For hyperbolic $\mathcal{C}U$, this is equivalent to the uniform quasiconvexity

from Definition 1.1, and is sufficient for our needs; for the present applications Definition 6.8 is sufficiently general, as well as for applications in [BHS17b].

Our goal is to prove the following two theorems, which provide hierarchy paths and a distance formula in relatively hierarchically hyperbolic spaces. We will not use these theorems in the remainder of this paper, but they are required for future applications.

Theorem 6.10 (Distance formula for relatively hierarchically hyperbolic spaces). *Let $(\mathcal{X}, \mathfrak{S})$ be a relatively hierarchically hyperbolic space. Then there exists s_0 such that for all $s \geq s_0$, there exist constants C, K such that for all $x, y \in \mathcal{X}$,*

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

Proof. By Proposition 6.15 below, for some suitably-chosen $\theta \geq 0$ and each $x, y \in \mathcal{X}$, there exists a subspace $M_\theta(x, y)$ of \mathcal{X} (endowed with the induced metric) so that $(M_\theta(x, y), \mathfrak{S})$ is a hierarchically hyperbolic space (with the same nesting relations and projections from $(\mathcal{X}, \mathfrak{S})$), so that for all $U \in \mathfrak{S}$, we have that $\pi_U(M_\theta(x, y)) \subset \mathcal{N}_\theta(\gamma_U)$, where γ_U is an arbitrarily-chosen geodesic in $\mathcal{C}U$ from $\pi_U(x)$ to $\pi_U(y)$. We emphasize that all of the constants from Definition 1.1 (for $M_\theta(x, y)$) are independent of x, y . The theorem now follows by applying the distance formula for hierarchically hyperbolic spaces (Theorem 4.5) to $(M_\theta(x, y), \mathfrak{S})$. \square

Theorem 6.11 (Hierarchy paths in relatively hierarchically hyperbolic spaces). *Let $(\mathcal{X}, \mathfrak{S})$ be a relatively hierarchically hyperbolic space. Then there exists $D \geq 0$ such that for all $x, y \in \mathcal{X}$, there exists a (D, D) -quasigeodesic γ in \mathcal{X} joining x, y so that $\pi_U(\gamma)$ is an unparameterized (D, D) -quasigeodesic.*

Proof. Proceed exactly as in Theorem 6.10, but apply Theorem 4.4 instead of Theorem 4.5. \square

We now define hulls of pairs of points in the relatively hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$. Let θ be a constant to be chosen (it will be the output of the realization theorem for a consistency constant depending on the constants associated to $(\mathcal{X}, \mathfrak{S})$), and let $x, y \in \mathcal{X}$. For each $U \in \mathfrak{S}$, fix a geodesic γ_U in $\mathcal{C}U$ joining $\pi_U(x)$ to $\pi_U(y)$. Define maps $r_U : \mathcal{C}U \rightarrow \gamma_U$ as follows: if $\mathcal{C}U$ is hyperbolic, let r_U be the coarse closest-point projection map. Otherwise, if $\mathcal{C}U$ is not hyperbolic (so U is \sqsubseteq -minimal), define r_U as follows: parametrize γ_U by arc length with $\gamma_U(0) = x$, and for each $p \in \mathcal{C}U$, let $m(p) = \min\{d_U(x, p), d_U(x, y)\}$. Then $r_U(p) = \gamma_U(m(p))$. This r_U is easily seen to be an L -coarsely Lipschitz retraction, with L independent of U and x, y . (When U is minimal, r_U is 1-Lipschitz.)

Next, define the hull $M_\theta(x, y)$ to be the set of points $x \in \mathcal{X}$ such that $d_U(x, \gamma_U) \leq \theta$ for all $U \in \mathfrak{S}$. In the next proposition, we show that $M_\theta(x, y)$ is a hierarchically hyperbolic space, with the following hierarchically hyperbolic structure:

- (1) the index set is \mathfrak{S} ;
- (2) the nesting, orthogonality, and transversality relations on \mathfrak{S} are the same as in $(\mathcal{X}, \mathfrak{S})$;
- (3) for each $U \in \mathfrak{S}$, the associated hyperbolic space is γ_U ;
- (4) for each $U \in \mathfrak{S}$, the projection $\pi'_U : M_\theta(x, y) \rightarrow \gamma_U$ is given by $\pi'_U = r_U \circ \pi_U$;
- (5) for each pair $U, V \in \mathfrak{S}$ of distinct non-orthogonal elements, the relative projection $\mathcal{C}U \rightarrow \mathcal{C}V$ is given by $r_V \circ \rho_V^U$.

Since there are now two sets of projections (those defined in the original hierarchical space $(\mathcal{X}, \mathfrak{S})$, denoted π_* , and the new projections π'_*), in the following proofs we will explicitly write all projections when writing distances in the various $\mathcal{C}U$.

Lemma 6.12 (Gates in hulls). *Let $M_\theta(x, y)$ be as above. Then there exists a uniformly coarsely Lipschitz retraction $r : \mathcal{X} \rightarrow M_\theta(x, y)$ such that for each $U \in \mathfrak{S}$, we have, up to uniformly (independent of x, y) bounded error, $\pi_U \circ r = r_U \circ \pi_U$.*

Remark 6.13. It is crucial in the following proof that \mathcal{CU} is δ -hyperbolic for each $U \in \mathfrak{S}$ that is not \sqsubseteq -minimal.

Proof of Lemma 6.12. Let $z \in \mathcal{X}$ and, for each U , let $t_U = r_U \circ \pi_U(z)$; this defines a tuple $(t_U) \in \prod_{U \in \mathfrak{S}} 2^{\mathcal{CU}}$ which we will check is κ -consistent for κ independent of x, y . The tuple (t_U) is admissible because of quasiconvexity of the images of projections to hyperbolic spaces, and coarse surjectivity of projections to non-hyperbolic ones.

Realization (Theorem 3.1) then yields $m \in \mathcal{X}$ such that $d_U(\pi_U(m), t_U) \leq \theta$ for all $U \in \mathfrak{S}$. By definition, $t_U \in \gamma_U$, so $m \in M_\theta(x, y)$ and we define $\mathfrak{g}_{x,y}(z) = m$. Note that up to perturbing slightly, we may take $\mathfrak{g}_{x,y}(z) = z$ when $z \in M_\theta$. Hence it suffices to check consistency of (t_U) .

First let $U, V \in \mathfrak{S}$ satisfy $U \triangleleft V$. Then $d_V(\pi_V(x), \pi_V(y)) \leq 2E$ (up to exchanging U and V), and moreover each of $\pi_U(x), \pi_U(y)$ is E -close to ρ_V^U . Since t_V lies on γ_V , it follows that $d_V(t_V, \rho_V^U) \leq 2E$.

Next, let $U, V \in \mathfrak{S}$ satisfy $U \not\sqsubseteq V$. Observe that in this case, \mathcal{CV} is δ -hyperbolic because V is not \sqsubseteq -minimal. First suppose that $d_V(\gamma_V, \rho_V^U) > 1$. Then by consistency and bounded geodesic image, $d_U(x, y) \leq 3E$, and $\text{diam}_U(\rho_V^U(\gamma_V)) \leq E$. It follows that $\text{diam}_U(t_U \cup \rho_V^U(t_V)) \leq 10E$.

Hence, suppose that $d_V(\rho_V^U, \gamma_V) \leq 10E$ but that $d_V(t_V, \rho_V^U) > E$. Without loss of generality, ρ_V^U lies at distance $\leq E$ from the subpath of γ_V joining t_V to $\pi_V(y)$. Let γ'_V be the subpath joining x to t_V . By consistency, bounded geodesic image, and the fact that \mathcal{CV} is δ -hyperbolic and $t_V = r_V \circ \pi_V(z)$, the geodesic triangle between $\pi_V(x), \pi_V(z)$, and t_V projects under ρ_V^U to a set, of diameter bounded by some uniform ξ , containing $\pi_U(x), \pi_U(z)$, and $\rho_V^U(t_V)$. Hence, since $t_U = r_U \circ \pi_U(z)$, and $\pi_U(x) \in \gamma_U$, the triangle inequality yields a uniform bound on $\text{diam}_U(t_U \cup \rho_V^U(t_V))$. Hence there exists a uniform κ , independent of x, y , so that (t_U) is κ -consistent. Finally, $\mathfrak{g}_{x,y}$ is coarsely Lipschitz by the uniqueness axiom (Definition 1.1.(9)), since each r_U is uniformly coarsely Lipschitz. \square

Lemma 6.14. *Let $m, m' \in M_\theta(x, y)$. Then there exists $C \geq 0$ such that m, m' are joined by a (C, C) -quasigeodesic in $M_\theta(x, y)$.*

Proof. Since \mathcal{X} is a quasigeodesic space, there exists $K \geq 0$ so that m, m' are joined by a K -discrete (K, K) -quasigeodesic $\sigma : [0, \ell] \rightarrow \mathcal{X}$ with $\sigma(0) = m, \sigma(\ell) = m'$. Note that $\mathfrak{g}_{x,y} \circ \sigma$ is a K' -discrete, efficient path for K' independent of x, y , since the gate map is uniformly coarsely Lipschitz. A minimal-length K' -discrete efficient path in $M_\theta(x, y)$ from x to y has the property that each subpath is K' -efficient, and is a uniform quasigeodesic, as needed. \square

Proposition 6.15. *For all sufficiently large θ , the data (1)–(5) above makes $(M_\theta(x, y), \mathfrak{S})$ a hierarchically hyperbolic space, where $M_\theta(x, y)$ inherits its metric as a subspace of \mathcal{X} . Moreover, the associated constants from Definition 1.1 are independent of x, y .*

Proof. By Lemma 6.14, $M_\theta(x, y)$ is a uniform quasigeodesic space. We now verify that the enumerated axioms from Definition 1.1 are satisfied. Each part of the definition involving only \mathfrak{S} and the $\sqsubseteq, \perp, \triangleleft$ relations is obviously satisfied; this includes finite complexity. The consistency inequalities hold because they hold in $(\mathcal{X}, \mathfrak{S})$ and each r_U is L -coarsely Lipschitz. The same holds for bounded geodesic image and the large link lemma. We now verify the two remaining claims:

Uniqueness: Let $m, m' \in M_\theta(x, y)$, so that $d_U(\pi_U(m), \gamma_U), d_U(\pi_U(m'), \gamma_U) \leq \theta$ for all $U \in \mathfrak{S}$. The definition of r_U ensures that $d_U(r_U \circ \pi_U(m), r_U \circ \pi_U(m')) \geq d_U(\pi_U(m), \pi_U(m')) - 2\theta$, and uniqueness follows.

Partial realization: Let $\{U_i\}$ be a totally orthogonal subset of \mathfrak{S} and choose, for each i , some $p_i \in \gamma_{U_i}$. By partial realization in $(\mathcal{X}, \mathfrak{S})$, there exists $z \in \mathcal{X}$ so that $d_{U_i}(\pi_{U_i}(z), p_i) \leq E$ for each i and $d_V(\pi_V(z), \rho_V^{U_i}) \leq E$ provided $U_i \sqsubset V$ or $U_i \pitchfork V$. Let $z' = \mathfrak{g}_{x,y}(z) \in M_\theta(x, y)$. Then, by the definition of the gate map and the fact that each r_U is L -coarsely Lipschitz, there exists α , independent of x, y , so that $d_{U_i}(r_{U_i} \circ \pi_{U_i}(z'), p_i) \leq \alpha$, while $d_V(r_V \circ \pi_V(z'), \rho_V^{U_i}) \leq \alpha$ whenever $U_i \pitchfork V$ or $U_i \sqsubset V$. Hence z' is the required partial realization point. This completes the proof that $(M_\theta(x, y), \mathfrak{S})$ is an HHS. \square

7. THE COARSE MEDIAN PROPERTY

In this section, we study the relationship between hierarchically hyperbolic spaces and spaces that are *coarse median* in the sense defined by Bowditch in [Bow13]. In particular, this discussion shows that $\text{Out}(F_n)$ is not a hierarchically hyperbolic space, and hence not a hierarchically hyperbolic group, for $n \geq 3$.

Definition 7.1 (Median graph). Let Γ be a graph with unit-length edges and path-metric d . Then Γ is a *median graph* if there is a map $\mathfrak{m}: \Gamma^3 \rightarrow \Gamma$ such that, for all $x, y, z \in \Gamma$, we have $d(x, y) = d(x, \mathfrak{m}(x, y, z)) + d(\mathfrak{m}(x, y, z), y)$, and likewise for the pairs x, z and y, z , where $\mathfrak{m}(x, y, z) = \mathfrak{m}(x, y, z)$. Note that if $x = y$, then $\mathfrak{m}(x, y, z) = x$.

Chepoi established in [Che00] that Γ is a median graph precisely when Γ is the 1-skeleton of a CAT(0) cube complex.

Definition 7.2 (Coarse median space). Let (M, d) be a metric space and let $\mathfrak{m}: M^3 \rightarrow M$ be a ternary operation satisfying the following:

- (1) (Triples) There exist constants $\kappa, h(0)$ such that for all $a, a', b, b', c, c' \in M$,

$$d(\mathfrak{m}(a, b, c), \mathfrak{m}(a', b', c')) \leq \kappa (d(a, a') + d(b, b') + d(c, c')) + h(0).$$

- (2) (Tuples) There is a function $h: \mathbb{N} \cup \{0\} \rightarrow [0, \infty)$ such that for any $A \subseteq M$ with $1 \leq |A| = p < \infty$, there is a CAT(0) cube complex \mathcal{F}_p and maps $\pi: A \rightarrow \mathcal{F}_p^{(0)}$ and $\lambda: \mathcal{F}_p^{(0)} \rightarrow M$ such that $d(a, \lambda(\pi(a))) \leq h(p)$ for all $a \in A$ and such that

$$d(\lambda(\mathfrak{m}_p(x, y, z)), \mathfrak{m}(\lambda(x), \lambda(y), \lambda(z))) \leq h(p)$$

for all $x, y, z \in \mathcal{F}_p$, where \mathfrak{m}_p is the map that sends triples from $\mathcal{F}_p^{(0)}$ to their median.

Then (M, d, \mathfrak{m}) is a *coarse median space*. The *rank* of (M, d, \mathfrak{m}) is at most d if each \mathcal{F}_p above can be chosen to satisfy $\dim \mathcal{F}_p \leq d$, and the rank of (M, d, \mathfrak{m}) is exactly d if d is the minimal integer such that (M, d, \mathfrak{m}) has rank at most d .

The next fact was observed by Bowditch [Bow18]; we include a proof for completeness.

Theorem 7.3 (Hierarchically hyperbolic implies coarse median). *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then \mathcal{X} is coarse median of rank at most the complexity of $(\mathcal{X}, \mathfrak{S})$.*

Proof. Since the spaces $\mathcal{C}U, U \in \mathfrak{S}$ are δ -hyperbolic for some δ independent of U , there exists for each U a ternary operation $\mathfrak{m}^U: \mathcal{C}U^3 \rightarrow \mathcal{C}U$ so that $(\mathcal{C}U, d_U, \mathfrak{m}^U)$ is a coarse median space of rank 1, and the constant κ and function $h: \mathbb{N} \cup \{0\} \rightarrow [0, \infty)$ from Definition 7.2 can be chosen to depend only on δ (and not on U).

Definition of the median: Define a map $\mathfrak{m}: \mathcal{X}^3 \rightarrow \mathcal{X}$ as follows. Let $x, y, z \in \mathcal{X}$ and, for each $U \in \mathfrak{S}$, let $b_U = \mathfrak{m}^U(\pi_U(x), \pi_U(y), \pi_U(z))$. By Lemma 2.6, the tuple $\vec{b} \in \prod_{U \in \mathfrak{S}} 2^{\mathcal{C}U}$ whose U -coordinate is b_U is κ -consistent for an appropriate choice of κ . Hence,

by the realization theorem (Theorem 3.1), there exists θ_e and $\mathbf{m} = \mathbf{m}(x, y, z) \in \mathcal{X}$ such that $d_U(\mathbf{m}, b_U) \leq \theta_u$ for all $U \in \mathfrak{S}$. Moreover, this is coarsely well-defined (up to the constant θ_e from the realization theorem).

Application of [Bow13, Proposition 10.1]: Observe that, by Definition 1.1.1, the projections $\pi_U: \mathcal{X} \rightarrow \mathcal{C}U$, $U \in \mathfrak{S}$ are uniformly coarsely Lipschitz. Moreover, for each $U \in \mathfrak{S}$, the projection $\pi_U: \mathcal{X} \rightarrow \mathcal{C}U$ is a “quasimorphism” in the sense of [Bow13, Section 10], i.e., $d_U(\mathbf{m}^U(\pi_U(x), \pi_U(y), \pi_U(z)), \pi_U(\mathbf{m}(x, y, z)))$ is uniformly bounded, by construction, as U varies over \mathfrak{S} and x, y, z vary in \mathcal{X} . Proposition 10.1 of [Bow13] then implies that \mathbf{m} is a coarse median on \mathcal{X} , since that the hypothesis (P1) of that proposition holds in our situation by the distance formula. \square

The following is a consequence of Theorem 7.3 and work of Bowditch [Bow13, Bow14a]:

Corollary 7.4 (Contractibility of asymptotic cones). *Let \mathcal{X} be a hierarchically hyperbolic space. Then all the asymptotic cones of \mathcal{X} are contractible, and in fact bi-Lipschitz equivalent to $CAT(0)$ spaces.*

Corollary 7.5 (HHGs have quadratic Dehn function). *Let G be a finitely generated group that is a hierarchically hyperbolic space. Then G is finitely presented and has quadratic Dehn function. In particular, this conclusion holds when G is a hierarchically hyperbolic group.*

Proof. This follows from Theorem 7.3 and [Bow13, Corollary 8.3]. \square

Corollary 7.6. *For $n \geq 3$, the group $\text{Out}(F_n)$ is not a hierarchically hyperbolic space, and in particular is not a hierarchically hyperbolic group.*

Proof. This is an immediate consequence of Corollary 7.5 and the exponential lower bound on the Dehn function of $\text{Out}(F_n)$ given by the combined results of [BV95, HM13b, BV]. \square

We also recover a special case of Theorem I of [BHS17a], using Corollary 7.5 and a theorem of Gersten-Holt-Riley [GHR03, Theorem A]:

Corollary 7.7. *Let N be a finitely generated virtually nilpotent group. Then G is quasi-isometric to a hierarchically hyperbolic space if and only if N is virtually abelian.*

Corollary 7.8. *Let S be a symmetric space of non-compact type, or a thick affine building. Suppose that the spherical type of S is not A_1^* . Then S is not hierarchically hyperbolic.*

Proof. This follows from Theorem 7.3 and Theorem A of [Hae16]. \square

Finally, Theorem 9.1 of [Bow14a] combines with Theorem 7.3 to yield:

Corollary 7.9 (Rapid decay). *Let G be a group whose Cayley graph is a hierarchically hyperbolic space. Then G has the rapid decay property.*

7.1. Coarse media and hierarchical quasiconvexity. The natural notion of quasiconvexity in the coarse median setting is related to hierarchical quasiconvexity.

Definition 7.10 (Coarsely convex). Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space and let $\mathbf{m}: \mathcal{X}^3 \rightarrow \mathcal{X}$ be the coarse median map constructed in the proof of Theorem 7.3. A closed subspace $\mathcal{Y} \subseteq \mathcal{X}$ is μ -convex if for all $y, y' \in \mathcal{Y}$ and $x \in \mathcal{X}$, we have $\mathbf{m}(y, y', x) \in \mathcal{N}_\mu(\mathcal{Y})$.

Remark 7.11. We will not use μ -convexity in the remainder of the paper. However, it is of independent interest since it parallels a characterization of convexity in median spaces: a subspace \mathcal{Y} of a median space is convex exactly when, for all $y, y' \in \mathcal{Y}$ and x in the ambient median space, the median of x, y, y' lies in \mathcal{Y} .

Proposition 7.12 (Coarse convexity and hierarchical quasiconvexity). *Let $(\mathcal{X}, \mathfrak{S})$ be a hierarchically hyperbolic space and let $\mathcal{Y} \subseteq \mathcal{X}$. If \mathcal{Y} is hierarchically k -quasiconvex, then there exists $\mu \geq 0$, depending only on k and the constants from Definition 1.1, such that \mathcal{Y} is μ -convex.*

Proof. Let $\mathcal{Y} \subseteq \mathcal{X}$ be k -hierarchically quasiconvex, let $y, y' \in \mathcal{Y}$ and $x \in \mathcal{X}$. Let $m = \mathbf{m}(x, y, y')$. For any $U \in \mathfrak{S}$, the projection $\pi_U(\mathcal{Y})$ is by definition $k(0)$ -quasiconvex, so that, for some $k' = k'(k(0), \delta)$, we have $d_U(m_U, \pi_U(\mathcal{Y})) \leq k'$, where m_U is the coarse median of $\pi_U(x), \pi_U(y), \pi_U(y')$ coming from hyperbolicity of $\mathcal{C}U$. The tuple $(m_U)_{U \in \mathfrak{S}}$ was shown above to be κ -consistent for appropriately-chosen κ (Lemma 2.6), and $d_U(m_U, \mathbf{m}(x, y, y')) \leq \theta_e(\kappa)$, so, by hierarchical quasiconvexity $d_{\mathcal{X}}(\mathbf{m}(x, y, y'), \mathcal{Y})$ is bounded by a constant depending on $k(\kappa)$ and k' . \square

8. COMBINATION THEOREMS FOR HIERARCHICALLY HYPERBOLIC SPACES

The goal of this section is to prove Theorem 8.6, which enables the construction of new hierarchically hyperbolic spaces and groups from a tree of given ones. We postpone the statement of the theorem until after the relevant definitions.

Definition 8.1 (Quasiconvex hieromorphism, full hieromorphism). Let $(f, f^\diamond, \{f^*(U)\}_{U \in \mathfrak{S}})$ be a hieromorphism $(\mathcal{X}, \mathfrak{S}) \rightarrow (\mathcal{X}', \mathfrak{S}')$. We say f is k -hierarchically quasiconvex if its image is k -hierarchically quasiconvex and $f: \mathcal{X} \rightarrow \mathcal{X}'$ is a quasi-isometric embedding. The hieromorphism is *full* if:

- (1) there exists $\xi \geq 0$ such that each $f^*(U): \mathcal{C}U \rightarrow \mathcal{C}(f^\diamond(U))$ is a (ξ, ξ) -quasi-isometry, and
- (2) for each $U \in \mathfrak{S}$, if $V' \in \mathfrak{S}'$ satisfies $V' \sqsubseteq f^\diamond(U)$, then there exists $V \in \mathfrak{S}$ such that $V \sqsubseteq U$ and $f^\diamond(V) = V'$.

Remark 8.2. Observe that Definition 8.1.(2) holds automatically unless V' is bounded.

Definition 8.3 (Tree of hierarchically hyperbolic spaces). Let \mathcal{V}, \mathcal{E} denote the vertex and edge-sets, respectively, of the simplicial tree T . A *tree of hierarchically hyperbolic spaces* is a quadruple

$$\mathcal{T} = (T, \{\mathcal{X}_v\}, \{\mathcal{X}_e\}, \{\phi_{e_\pm} : v \in \mathcal{V}, e \in \mathcal{E}\})$$

satisfying:

- (1) $\{\mathcal{X}_v\}$ and $\{\mathcal{X}_e\}$ are uniformly hierarchically hyperbolic: each \mathcal{X}_v has index set \mathfrak{S}_v , and each \mathcal{X}_e has index set \mathfrak{S}_e . In particular, there is a uniform bound on the complexities of the hierarchically hyperbolic structures on the \mathcal{X}_v and \mathcal{X}_e .
- (2) Fix an orientation on each $e \in \mathcal{E}$ and let e_+, e_- denote the initial and terminal vertices of e . Then, each $\phi_{e_\pm}: \mathcal{X}_e \rightarrow \mathcal{X}_{e_\pm}$ is a hieromorphism with all constants bounded by some uniform $\xi \geq 0$. (We adopt the hieromorphism notation from Definition 1.20. Hence we actually have maps $\phi_{e_\pm}: \mathcal{X}_e \rightarrow \mathcal{X}_{e_\pm}$, and maps $\phi_{e_\pm}^\diamond: \mathfrak{S}_e \rightarrow \mathfrak{S}_{e_\pm}$ preserving nesting, transversality, and orthogonality, and coarse ξ -Lipschitz maps $\phi_{e_\pm}^*(U): \mathcal{C}U \rightarrow \mathcal{C}(\phi_{e_\pm}^\diamond(U))$ satisfying the conditions of that definition.)

Given a tree \mathcal{T} of hierarchically hyperbolic spaces, denote by $\mathcal{X}(\mathcal{T})$ the metric space constructed from $\bigsqcup_{v \in \mathcal{V}} \mathcal{X}_v$ by adding edges of length 1 as follows: if $x \in \mathcal{X}_e$, we declare $\phi_{e_-}(x)$ to be joined by an edge to $\phi_{e_+}(x)$. Given $x, x' \in \mathcal{X}$ in the same vertex space \mathcal{X}_v , define $d'(x, x')$ to be $d_{\mathcal{X}_v}(x, x')$. Given $x, x' \in \mathcal{X}$ joined by an edge, define $d'(x, x') = 1$. Given a sequence $x_0, x_1, \dots, x_k \in \mathcal{X}$, with consecutive points either joined by an edge or in a common vertex space, define its length to be $\sum_{i=1}^{k-1} d'(x_i, x_{i+1})$. Given $x, x' \in \mathcal{X}$, let $d(x, x')$ be the infimum of the lengths of such sequences $x = x_0, \dots, x_k = x'$.

Remark 8.4. Since the vertex spaces are (uniform) quasigeodesic spaces, (\mathcal{X}, d) is a quasigeodesic space.

Definition 8.5 (Equivalence, support, bounded support). Let \mathcal{T} be a tree of hierarchically hyperbolic spaces. For each $e \in \mathcal{E}$, and each $W_{e_-} \in \mathfrak{S}_{e_-}, W_{e_+} \in \mathfrak{S}_{e_+}$, write $W_{e_-} \sim_d W_{e_+}$ if there exists $W_e \in \mathfrak{S}_e$ so that $\phi_{e_\pm}^\diamond(W_e) = W_{e_\pm}$. The transitive closure \sim of \sim_d is an equivalence relation on $\bigcup_v \mathfrak{S}_v$. The \sim -class of $W \in \bigcup_v \mathfrak{S}_v$ is denoted $[W]$.

The *support* of an equivalence class $[W]$ is the induced subgraph $T_{[W]}$ of T whose vertices are those $v \in T$ so that \mathfrak{S}_v contains a representative of $[W]$. Observe that $T_{[W]}$ is connected. The tree \mathcal{T} of hierarchically hyperbolic spaces has *bounded supports* if there exists $n \in \mathbb{N}$ such that each \sim -class has support of diameter at most n .

We can now state the main theorem of this section:

Theorem 8.6 (Combination theorem for hierarchically hyperbolic spaces). *Let \mathcal{T} be a tree of hierarchically hyperbolic spaces. Suppose that:*

- (1) *there exists a function k so that each edge-hieromorphism is k -hierarchically quasi-conver;*
- (2) *each edge-hieromorphism is full;*
- (3) *\mathcal{T} has bounded supports of diameter at most n ;*
- (4) *if e is an edge of T and S_e is the \sqsubseteq -maximal element of \mathfrak{S}_e , then for all $V \in \mathfrak{S}_{e_\pm}$, the elements V and $\phi_{e_\pm}^\diamond(S_e)$ are not orthogonal in \mathfrak{S}_{e_\pm} . Moreover, there exists $K \geq 0$ such that for all vertices v of \mathcal{T} and edges e incident to v , we have $d_{\text{Haus}}(\phi_v(\mathcal{X}_e), \mathbf{F}_{\phi_v^\diamond(S_e)} \times \{\star\}) \leq K$, where $S_e \in \mathfrak{S}_e$ is the unique maximal element and $\star \in \mathbf{E}_{\phi_v^\diamond(S_e)}$.*

Then $\mathcal{X}(\mathcal{T})$ is hierarchically hyperbolic.

We postpone the proof until after the necessary lemmas and definitions. For the remainder of this section, fix a tree of hierarchically hyperbolic spaces $\mathcal{T} = (T, \{\mathcal{X}_v\}, \{\mathcal{X}_e\}, \{\phi_{e_\pm}\})$ satisfying the hypotheses of Theorem 8.6; let n be the constant implicit in assumption 3.

Let $\mathfrak{S}^0 = \{T\} \cup (\bigcup_v \mathfrak{S}_v / \sim)$.

Definition 8.7 (Nesting, orthogonality, transversality in \mathfrak{S}^0). For all $[W] \in \mathfrak{S}$, declare $[W] \sqsubseteq T$. If $[V], [W]$ are \sim -classes, then $[V] \sqsubseteq [W]$ if and only if there exists $v \in T$ such that $[V], [W]$ are respectively represented by $V_v, W_v \in \mathfrak{S}_v$ and $V_v \sqsubseteq W_v$; this relation is *nesting*. For convenience, for $A \in \mathfrak{S}$, we write \mathfrak{S}_A to denote the set of $B \in \mathfrak{S}^0$ such that $B \sqsubseteq A$.

Likewise, $[V] \perp [W]$ if and only if there exists a vertex $v \in T$ such that $[V], [W]$ are respectively represented by $V_v, W_v \in \mathfrak{S}_v$ and $V_v \perp W_v$; this relation is *orthogonality*. If $[V], [W] \in \mathfrak{S}$ are not orthogonal and neither is nested into the other, then they are *transverse*, written $[V] \pitchfork [W]$. Equivalently, $[V] \pitchfork [W]$ if for all $v \in T_{[V]} \cap T_{[W]}$, the representatives $V_v, W_v \in \mathfrak{S}_v$ of $[V], [W]$ satisfy $V_v \pitchfork W_v$.

Fullness (Definition 8.1.(2)) was introduced to enable the following two lemmas:

Lemma 8.8. *Let \mathcal{T} be a tree of hierarchically hyperbolic spaces, let v be a vertex of the underlying tree T , and let $U, U' \in \mathfrak{S}_v$ satisfy $U \sqsubseteq U'$. Then either $U = U'$ or $U \not\sim U'$.*

Proof. Suppose that $U \sim U'$, so that there is a closed path $v = v_0, v_1, \dots, v_n = v$ in T and a sequence $U = U_0, U_1, \dots, U_n = U'$ such that $U_i \in \mathfrak{S}_{v_i}$ and $U_i \sim_d U_{i+1}$ for all i . If $U \neq U'$, then Condition (2) (fullness) from Definition 8.1 and the fact that hieromorphisms preserve nesting yields $U'' \in \mathfrak{S}_v$, different from U' , such that $U'' \sim U$ and $U'' \sqsubset U' \sqsubset U$ (where \sqsubset denotes proper nesting). Repeating this argument contradicts finiteness of complexity. \square

Lemma 8.9. *The relation \sqsubseteq is a partial order on \mathfrak{S}^0 , and T is the unique \sqsubseteq -maximal element. Moreover, if $[V] \perp [W]$ and $[U] \sqsubseteq [V]$, then $[U] \perp [W]$ and $[V], [W]$ are not \sqsubseteq -comparable.*

Proof. Reflexivity is clear. Suppose that $[V_v] \sqsubseteq [U_u] \sqsubseteq [W_w]$. Then there are vertices $v_1, v_2 \in \mathcal{V}$ and representatives $V_{v_1} \in [V_v], U_{v_1} \in [U_u], U_{v_2} \in [U_u], W_{v_2} \in [W_w]$ so that $V_{v_1} \sqsubseteq U_{v_1}$ and $U_{v_2} \sqsubseteq W_{v_2}$. Since edge-hieromorphisms are full, induction on $d_T(v_1, v_2)$ yields $V_{v_2} \sqsubseteq U_{v_2}$ so that $V_{v_2} \sim V_{v_1}$. Transitivity of the nesting relation in \mathfrak{S}_{v_2} implies that $V_{v_2} \sqsubseteq W_{v_2}$, whence $[V_v] \sqsubseteq [W_w]$.

Suppose that $[U_u] \sqsubseteq [V_v]$ and $[V_v] \sqsubseteq [U_u]$, and suppose by contradiction that $[U_u] \neq [V_v]$. Choose $v_1, v_2 \in \mathcal{V}$ and representatives $U_{v_1}, U_{v_2}, V_{v_1}, V_{v_2}$ so that $U_{v_1} \sqsubseteq V_{v_1}$ and $V_{v_2} \sqsubseteq U_{v_2}$. The definition of \sim again yields $U_{v_2} \sim U_{v_1}$ with $U_{v_2} \sqsubseteq V_{v_2} \neq U_{v_2}$. This contradicts Lemma 8.8. Hence \sqsubseteq is antisymmetric, whence it is a partial order. The underlying tree T is the unique \sqsubseteq -maximal element by definition.

Suppose that $[V] \perp [W]$ and $[U] \sqsubseteq [V]$. Then there are vertices v_1, v_2 and representatives $V_{v_1}, W_{v_1}, U_{v_2}, V_{v_2}$ such that $V_{v_1} \perp W_{v_1}$ and $U_{v_2} \perp V_{v_2}$. Again by fullness of the edge-hieromorphisms, there exists $U_{v_1} \sim U_{v_2}$ with $U_{v_1} \sqsubseteq V_{v_1}$, whence $U_{v_1} \perp W_{v_1}$. Thus $[U] \perp [W]$ as required. Also, \sqsubseteq -incomparability of $[V], [W]$ follows from fullness and the fact that edge-hieromorphisms preserve orthogonality and nesting. \square

Lemma 8.10. *Let $[W] \in \mathfrak{S}^0$ and let $[U] \sqsubseteq [W]$. Suppose moreover that*

$$\{[V] \in \mathfrak{S}_{[W]} : [V] \perp [U]\} \neq \emptyset.$$

Then there exists $[A] \in \mathfrak{S}_{[W]} - \{[W]\}$ such that $[V] \sqsubseteq [A]$ for all $[V] \in \mathfrak{S}_{[W]}$ with $[V] \perp [U]$.

Proof. Choose some $v \in \mathcal{V}$ so that there exists $V_v \in \mathfrak{S}_v$ and $U_v \in \mathfrak{S}_v$ with $[U_v] = [U]$ and $V_v \perp U_v$. Then by definition, there exists $A_v \in \mathfrak{S}_v$ so that $B_v \sqsubseteq A_v$ whenever $B_v \perp U_v$ and so that $[B_v] \sqsubseteq [W]$. It follows from the fact that the edge-hieromorphisms are full and preserve (non)orthogonality that $[B] \sqsubseteq [A_v]$ whenever $[B] \perp [U]$. \square

The set \mathfrak{S}^0 is not quite large enough to satisfy the orthogonality axiom, for the following reason: in Lemma 8.10, we needed $[W]$ to be a \sim -class, but since $T \in \mathfrak{S}^0$, we need to be able to satisfy the axiom with $[W]$ replaced by T . To this end, we add some new elements to \mathfrak{S}^0 , and extend the $\sqsubseteq, \perp, \triangleleft$ relations, as follows.

Definition 8.11 (Containers and \mathfrak{S}). We now define the index set \mathfrak{S} for the HHS structure we will construct in order to prove Theorem 8.6. First, \mathfrak{S} contains \mathfrak{S}^0 . Next, for each $[W]$ for which there exists $[U]$ with $[U] \perp [W]$, let $K_0^\perp([W])$ be a new element of \mathfrak{S} , which we call the *container* of $[W]$. We make the following declarations:

- $K_0^\perp([W]) \sqsubseteq T$;
- $[U] \sqsubseteq K_0^\perp([W])$ if and only if $[U] \perp [W]$;
- $K_0^\perp([W]) \triangleleft K_0^\perp([U])$ if $[U] \neq [W]$;
- $K_0^\perp([W]) \perp [V]$ if and only if $[V] \sqsubseteq [W]$;
- for all other $[U]$, we have $[U] \triangleleft K_0^\perp([W])$.

Let \mathcal{K}_0 be the set of all $K_0^\perp([W])$ as $[W]$ varies among those \sim -classes for which there is at least one orthogonal \sim -class.

Next, for each $K_0^\perp([W]) \in \mathcal{K}_0$, consider a \sim -class $[U] \sqsubseteq K_0^\perp([W])$ such that $[U] \perp [V]$ for some other $[V] \sqsubseteq K_0^\perp([W])$. Let $K_1^\perp([W], [U])$ be a new element of \mathfrak{S} , and let \mathcal{K}_1 be the set of such containers, as $[W]$ varies and as $[U]$ varies over those \sim -classes nested in $K_0^\perp([W])$ (i.e. orthogonal to $[W]$) that are orthogonal to some other \sim -class nested in $K_0^\perp([W])$.

We now make the following declarations:

- $K_1^\perp([W], [U]) \sqsubseteq K_0^\perp([W])$ and $K_1^\perp([W], [U])$ is transverse to every other element of $\mathcal{K}_0 \cup \mathcal{K}_1$.
- $K_1^\perp([W], [U]) \sqsubseteq T$.
- $[V] \sqsubseteq K_1^\perp([W], [U])$ if and only if $[V] \sqsubseteq K_0^\perp([W])$ and $[V] \perp [U]$.
- $[V] \perp K_1^\perp([W], [U])$ if and only if either $[V] \sqsubseteq [W]$ (i.e. $[V] \perp K_0^\perp([W])$) or $[V] \sqsubseteq [U]$.
- If none of the two preceding conditions is satisfied by $[V]$, then $[V] \pitchfork K_1^\perp([W], [U])$.

We now proceed as above to inductively construct sets $\mathcal{K}_\eta, \eta \geq 1$ of new “containers”, where each $K_\eta^\perp([W]_1, \dots, [W_\eta])$ is nested in $K_i^\perp([W]_1, \dots, [W_i])$ for $i \leq \eta - 1$, and also nested in T . Our inductive construction ensures that $[W_1], \dots, [W_\eta]$ are pairwise-orthogonal. The \sim -classes $[U]$ nested in $K_\eta^\perp([W]_1, \dots, [W_\eta])$ are precisely those that are orthogonal to each of $[W_1], \dots, [W_\eta]$. The \sim -classes U orthogonal to $K_\eta^\perp([W]_1, \dots, [W_\eta])$ are precisely those $[U]$ nested into some $[W_i]$.

Let $\mathfrak{S} = \mathfrak{S}^0 \cup \bigcup_{\eta \geq 0} \mathcal{K}_\eta$.

Remark 8.12 (Extension of $\sqsubseteq, \perp, \pitchfork$ satisfies the axioms). Lemma 8.9 shows that \sqsubseteq is a partial order on \mathfrak{S}^0 , and Definition 8.11 shows how to extend \sqsubseteq to all of \mathfrak{S} . By construction, the extended \sqsubseteq continues to be transitive. This follows from Lemma 8.9, the definition, and induction on the η in \mathcal{K}_η . By definition, T is still the unique \sqsubseteq -maximal element.

Now suppose that $[U] \sqsubseteq K_\eta^\perp([W]_1, \dots, [W_\eta])$ and $[V] \perp K_\eta^\perp([W]_1, \dots, [W_\eta])$. Then $[V] \sqsubseteq [W_i]$ for some i , and $[U] \perp [W_j]$ for all j . Lemma 8.9 implies $[U] \perp [V]$. On the other hand, $K_\eta^\perp([W]_1, \dots, [W_\eta])$ is never nested into any \sim -class or orthogonal to any element of $\bigcup_{\eta} \mathcal{K}_\eta$.

Lemma 8.13. *There exists $\chi \geq 0$ such that if $\{V_1, \dots, V_c\} \subset \mathfrak{S}$ consists of pairwise orthogonal or pairwise \sqsubseteq -comparable elements, then $c \leq \chi$. In particular, $\bigcup_{\eta \geq 0} \mathcal{K}_\eta = \bigcup_{\eta=0}^{(\chi-1)/2} \mathcal{K}_\eta$.*

Proof. For each $v \in T$, let χ_v be the complexity of $(\mathcal{X}_v, \mathfrak{S}_v)$ and let $\chi = 2 \max_v \chi_v + 1$. Let $[V_1], \dots, [V_c] \in \mathfrak{S} - \{T\}$ be \sim -classes that are pairwise orthogonal or pairwise \sqsubseteq -comparable. The Helly property for trees yields a vertex v lying in the support of each $[V_i]$; let $V_i^v \in \mathfrak{S}_v$ represent $[V_i]$. Since edge-hieromorphisms preserve nesting, orthogonality, and transversality, $c \leq \chi_v$.

Any pairwise-orthogonal set in \mathfrak{S} either has cardinality ≤ 1 or contains at most one element that is not a \sim -class, so the bound on pairwise-orthogonal sets is $\max_v \chi_v + 1$.

Hence it suffices to bound \sqsubseteq -chains in \mathfrak{S} . Any \sqsubseteq -chain $V_1 \sqsubseteq V_2 \sqsubseteq \dots \sqsubseteq V_k$ has the property that, for some $0 \leq m \leq k$, the first m elements are \sim -classes, and the remaining elements lie in $\{T\} \cup \bigcup_{\eta} \mathcal{K}_\eta$. Hence it suffices to show that any \sqsubseteq -chain in $\bigcup_{\eta \geq 0} \mathcal{K}_\eta$ has length at most χ . But by definition, any such chain has the form

$$K_0^\perp([W_0]) \supset K_1^\perp([W_0], [W_1]) \supset \dots \supset K_\eta^\perp([W_0], \dots, [W_{\eta-1}]),$$

where the $[W_i]$ are pairwise orthogonal. Hence $\eta \leq \max_v \chi_v \leq (\chi - 1)/2$, as required. This also proves the final assertion. \square

Definition 8.14 (Favorite representative, hyperbolic spaces associated to elements of \mathfrak{S}). Let $\mathcal{CT} = T$. For each \sim -class $[W]$, choose a *favorite vertex* v of $T_{[W]}$ and let $W_v \in \mathfrak{S}_{W_v}$ be the *favorite representative* of $[W]$. Let $\mathcal{C}[W] = \mathcal{C}W_v$. Note that each $\mathcal{C}[W]$ is δ -hyperbolic, where δ is the uniform hyperbolicity constant for \mathcal{T} .

Finally, for each $K^\perp \in \bigcup_{\eta} \mathcal{K}_\eta$, let $\mathcal{C}K^\perp$ be a single point.

Definition 8.15 (Gates in vertex spaces). For each vertex v of T , define a *gate map* $\mathfrak{g}_v: \mathcal{X} \rightarrow \mathcal{X}_v$ as follows. Let $x \in \mathcal{X}_u$ for some vertex u of T . We define $\mathfrak{g}_v(x)$ inductively on $d_T(u, v)$. If $u = v$, then set $\mathfrak{g}_v(x) = x$. Otherwise, $u = e^-$ for some edge e of T so that $d_T(e^+, v) = d_T(u, v) - 1$. Then set $\mathfrak{g}_v(x) = \mathfrak{g}_v(\phi_{e^+}(\phi_{e^-}^{-1}(\mathfrak{g}_{\phi_{e^-}(\mathcal{X}_e)}(x))))$. We also have a map $\beta_{v_e}: \mathcal{X} \rightarrow$

\mathcal{CV}_v , defined by $\beta_{V_v}(x) = \pi_{V_v}(\mathfrak{g}_v(x))$. (Here, $\mathfrak{g}_{\phi_{e^-}(\mathcal{X}_e)}: \mathcal{X}_{e^-} = \mathcal{X}_u \rightarrow \phi_{e^-}(\mathcal{X}_e)$ is the usual gate map to a hierarchically quasiconvex subspace, described in Definition 5.4, and $\phi_{e^\pm}^{-1}$ is a quasi-inverse for the edge-hieromorphism.)

Lemma 8.16. *There exists K , depending only on E and ξ , such that the following holds. Let e, f be edges of T and v a vertex so that $e^- = f^- = v$. Suppose for some $V \in \mathfrak{S}_v$ that there exist $x, y \in \phi_{e^-}(\mathcal{X}_e) \subseteq \mathcal{X}_v$ with $d_V(\mathfrak{g}_{\phi_{f^-}(\mathcal{X}_f)}(x), \mathfrak{g}_{\phi_{f^-}(\mathcal{X}_f)}(y)) > 10K$. Then $V \in \phi_{e^-}^\diamond(\mathfrak{S}_e) \cap \phi_{f^-}^\diamond(\mathfrak{S}_f)$.*

Proof. Let $\mathcal{Y}_e = \phi_{e^-}(\mathcal{X}_e)$ and let $\mathcal{Y}_f = \phi_{f^-}(\mathcal{X}_f)$; these spaces are uniformly hierarchically quasiconvex in \mathcal{X}_v . Moreover, by fullness of the edge-hieromorphisms, we can choose $K \geq 100E$ so that the map $\pi_V: \mathcal{Y}_e \rightarrow \mathcal{CV}$ is K -coarsely surjective for each $V \in \phi_{e^-}^\diamond(\mathfrak{S}_e)$, and likewise for $\phi_{f^-}^\diamond(\mathfrak{S}_f)$ and \mathcal{Y}_v . If $V \in \mathfrak{S}_v - \phi_{f^-}^\diamond(\mathfrak{S}_f)$, then π_V is K -coarsely constant on \mathcal{Y}_f , by the distance formula, since \mathcal{X}_f is quasi-isometrically embedded. Likewise, π_V is coarsely constant on \mathcal{Y}_e if $V \notin \phi_{e^-}^\diamond(\mathfrak{S}_e)$. (This also follows from consistency when V is transverse to some unbounded element of $\phi_{e^-}^\diamond(\mathfrak{S}_e)$ and from consistency and bounded geodesic image otherwise.)

Suppose that there exists $V \in \mathfrak{S}_v$ such that $d_V(\mathfrak{g}_{\phi_{f^-}(\mathcal{X}_f)}(x), \mathfrak{g}_{\phi_{f^-}(\mathcal{X}_f)}(y)) > 10K$. Since $\mathfrak{g}_{\phi_{f^-}(\mathcal{X}_f)}(x), \mathfrak{g}_{\phi_{f^-}(\mathcal{X}_f)}(y) \in \mathcal{X}_f$, we therefore have that $V \in \phi_{f^-}^\diamond(\mathfrak{S}_f)$. On the other hand, the definition of gates implies that $d_V(x, y) > 8K$, so $V \in \phi_{e^-}^\diamond(\mathfrak{S}_e)$. \square

Lemma 8.17. *There exists a constant K' such that the following holds. Let e, f be edges of T and suppose that there do not exist $V_e \in \mathfrak{S}_e, V_f \in \mathfrak{S}_f$ for which $\phi_{e^-}^\diamond(V_e) \sim \phi_{f^-}^\diamond(V_f)$. Then $\mathfrak{g}_{e^-}(\mathcal{X}_f)$ has diameter at most K' . In particular, the conclusion holds if $d_T(e, f) > n$, where n bounds the diameter of the supports.*

Proof. The second assertion follows immediately from the first in light of how n was chosen.

We now prove the first assertion by induction on the number k of vertices on the geodesic in T from e to f . The base case, $k = 1$, follows from Lemma 8.16.

For $k \geq 1$, let v_0, v_1, \dots, v_k be the vertices on a geodesic from e to f , in the obvious order. Let b be the edge joining v_{k-1} to v_k , with $b^- = v_k$.




FIGURE 4. Schematic of the subset of \mathcal{X} near \mathcal{X}_{b^-} .

It follows from the definition of gates that $\mathfrak{g}_{e^-}(\mathcal{X}_f)$ has diameter (coarsely) bounded above by that of $\mathfrak{g}_{\phi_{b^-}(\mathcal{X}_b)}(\mathcal{X}_f)$ and that of $\mathfrak{g}_{e^-}(\mathcal{X}_b)$. Hence suppose that $\text{diam}(\mathfrak{g}_{\phi_{b^-}(\mathcal{X}_b)}(\mathcal{X}_f)) > 10K$ and $\text{diam}(\mathfrak{g}_{e^-}(\mathcal{X}_b)) > 10K$. Then, by induction and Lemma 8.16, we see that there exists $V_e \in \mathfrak{S}_e, V_f \in \mathfrak{S}_f$ for which $\phi_{e^-}^\diamond(V_e) \sim \phi_{f^-}^\diamond(V_f)$, a contradiction. \square

Lemma 8.18. *The map $\mathfrak{g}_v: \mathcal{X} \rightarrow \mathcal{X}_v$ is coarsely Lipschitz, with constants independent of v .*

Proof. Let $x, y \in \mathcal{X}$. If the projections of x, y to T lie in the ball of radius $2n + 1$ about v , then this follows since \mathfrak{g}_v is the composition of a bounded number of maps, each of which

is uniformly coarsely Lipschitz by Lemma 5.5. Otherwise, by Remark 8.4, it suffices to consider x, y with $d_{\mathcal{X}}(x, y) \leq C$, where C depends only on the metric d . In this case, let v_x, v_y be the vertices in T to which x, y project. Let v' be the median in T of v, v_x, v_y . Observe that there is a uniform bound on $d_T(v_x, v')$ and $d_T(v_y, v')$, so it suffices to bound $d_v(\mathfrak{g}_v(\mathfrak{g}_{v'}(x)), \mathfrak{g}_v(\mathfrak{g}_{v'}(y)))$. Either $d_T(v, v') \leq 2n + 1$, and we are done, or Lemma 8.17 gives the desired bound, since equivalence classes have support of diameter at most n . \square

Definition 8.19 (Projections). For each $[W] \in \mathfrak{S}$, define the *projection* $\pi_{[W]}: \mathcal{X} \rightarrow \mathcal{C}[W]$ by $\pi_{[W]}(x) = \beta_{W_v}(x)$, where W_v is the favorite representative of $[W]$. Note that these projections take points to uniformly bounded sets, since the collection of vertex spaces is uniformly hierarchically hyperbolic. Define $\pi_T: \mathcal{X} \rightarrow T$ to be the usual projection to T . Finally, for each $K^\perp \in \bigcup_\eta \mathcal{K}_\eta$, just let $\pi_{K^\perp}: \mathcal{X} \rightarrow \mathcal{C}K^\perp$ be a constant map.

Lemma 8.20 (Comparison maps). *There exists a uniform constant $\xi \geq 1$ such that for all $W_v \in \mathfrak{S}_v, W_w \in \mathfrak{S}_w$ with $W_v \sim W_w$, there exists a (ξ, ξ) -quasi-isometry $\mathfrak{c}: W_v \rightarrow W_w$ such that $\mathfrak{c} \circ \beta_v = \beta_w$ up to uniformly bounded error.*

Definition 8.21. A map \mathfrak{c} as given by Lemma 8.20 is called a *comparison map*.

Proof of Lemma 8.20. We first clarify the situation by stating some consequences of the definitions. Let e^+, e^- be vertices of T joined by an edge e . Suppose that there exists $W^+ \in \mathfrak{S}_{e^+}, W^- \in \mathfrak{S}_{e^-}$ such that $W^+ \sim W^-$, so that there exists $W \in \mathfrak{S}_e$ with $\pi(\phi_{e^\pm})(W) = W^\pm$. Then the following diagram coarsely commutes (with uniform constants):

$$\begin{array}{ccccc}
 & & \mathcal{X} & & \\
 & \swarrow \mathfrak{g}_{e^-} & & \searrow \mathfrak{g}_{e^+} & \\
 & \mathcal{X}_{e^-} & & \mathcal{X}_{e^+} & \\
 \swarrow \pi_{W^-} & & & & \searrow \pi_{W^-} \\
 & \mathcal{C}W^- & \mathcal{X}_e & \mathcal{C}W^+ & \\
 & \swarrow & \downarrow \pi_W & \searrow & \\
 & & \mathcal{C}W & &
 \end{array}$$

where $\mathcal{X}_e \rightarrow \mathcal{X}_{e^\pm}$ is the uniform quasi-isometry ϕ_{e^\pm} , while $\mathcal{X}_{e^\pm} \rightarrow \mathcal{X}_e$ is the composition of a quasi-inverse for ϕ_{e^\pm} with the gate map $\mathcal{X}_{e^\pm} \rightarrow \phi_{e^\pm}(\mathcal{X}_e)$, and the maps $\mathcal{C}W \leftrightarrow \mathcal{C}W^\pm$ are the quasi-isometries implicit in the edge hieromorphism or their quasi-inverses. The proof essentially amounts to chaining together a sequence of these diagrams as e varies among the edges of a geodesic from v to w ; an important ingredient is played by the fact that such a geodesic has length at most n .

Let $v = v_0, v_1, \dots, v_m, v_{m+1} = w$ be the geodesic sequence in T from v to w and let e^i be the edge joining v_i to v_{i+1} . For each i , choose $W_i \in \mathfrak{S}_{e^i}$ and $W_i^\pm \in \mathfrak{S}_{e_i^\pm}$ so that (say) $W_0^- = W_v$ and $W_m^+ = W_w$ and so that $\phi_{e_i^\pm}^\diamond(W_i) = W_i^\pm$ for all i . For each i , let $q_i^\pm: \mathcal{C}W_i \rightarrow \mathcal{C}W_i^\pm$ be $q_i = \phi_{e_i^\pm}^*(W_i)$ be the (ξ', ξ') -quasi-isometry packaged in the edge-hieromorphism, and let \bar{q}_i^\pm be a quasi-inverse; the constant ξ' is uniform by hypothesis, and $m \leq n$ since \mathcal{T} has bounded supports. The hypotheses on the edge-hieromorphisms ensure that the W_i^\pm are uniquely determined by W_v, W_w , and we define \mathfrak{c} by

$$\mathfrak{c} = q_m^{\epsilon_m} \bar{q}_m^{\epsilon'_m} \dots q_1^{\epsilon_1} \bar{q}_1^{\epsilon'_1},$$

where $\epsilon_i, \epsilon'_i \in \{\pm\}$ depend on the orientation of e^i , and $\epsilon'_i = +$ if and only if $\epsilon_i = -$. This is a (ξ, ξ) -quasi-isometry, where $\xi = \xi(\xi'_n)$.

If $v = w$, then \mathbf{c} is the identity and $\mathbf{c} \circ \beta_v = \beta_v$. Let $d \geq 1 = d_T(v, w)$ and let w' be the penultimate vertex on the geodesic of T from v to w . Let $\mathbf{c}' : \mathcal{C}W_v \rightarrow \mathcal{C}W_{w'}$ be a comparison map, so that, by induction, there exists $\lambda' \geq 0$ so that $d_{\mathcal{C}W_{w'}}(\mathbf{c}' \circ \beta_v(x), \beta_{w'}(x)) \leq \lambda'$ for all $x \in \mathcal{X}$. Let $\mathbf{c}'' = \bar{q}_k^+ q_k^- : \mathcal{C}W_{w'} \rightarrow \mathcal{C}W_w$ be the (ξ', ξ') -quasi-isometry packaged in the edge-hieromorphism, so that the following diagram coarsely commutes:

$$\begin{array}{ccccc}
 & & \mathcal{X} & & \\
 & \swarrow \mathfrak{g}_v & \downarrow \mathfrak{g}_{w'} & \searrow \mathfrak{g}_w & \\
 \mathcal{X}_v & \xrightarrow{\mathfrak{g}_{w'}} & \mathcal{X}_{w'} & \xrightarrow{\mathfrak{g}_w} & \mathcal{X}_w \\
 \pi_{W_v} \downarrow & & \pi_{W_{w'}} \downarrow & & \pi_{W_w} \downarrow \\
 \mathcal{C}W_v & \xrightarrow{\mathbf{c}'} & \mathcal{C}W_{w'} & \xrightarrow{\mathbf{c}''} & \mathcal{C}W_w
 \end{array}$$

Since $\mathbf{c} = \mathbf{c}'' \circ \mathbf{c}'$ and the constants implicit in the coarse commutativity of the diagram depend only on the constants of the hieromorphism and on $d \leq n$, the claim follows. \square

Lemma 8.22. *There exists K such that each $\pi_{[W]}$ is (K, K) -coarsely Lipschitz.*

Proof. For each vertex v of T and each $V \in \mathfrak{S}_v$, the projection $\pi_V : \mathcal{X}_v \rightarrow \mathcal{C}V$ is uniformly coarsely Lipschitz, by definition. By Lemma 8.18, each gate map $\mathfrak{g}_v : \mathcal{X} \rightarrow \mathcal{X}_v$ is uniformly coarsely Lipschitz. The lemma follows since $\pi_{[W]} = \pi_{W_v} \circ \mathfrak{g}_v$, where v is the favorite vertex carrying the favorite representative W_v of $[W]$. \square

Definition 8.23 (Projections between hyperbolic spaces). If $[V] \sqsubseteq [W]$, then choose vertices $v, v', w \in \mathcal{V}$ so that V_v, W_w are respectively the favorite representatives of $[V], [W]$, while $V_{v'}, W_{v'}$ are respectively representatives of $[V], [W]$ with $V_{v'}, W_{v'} \in \mathfrak{S}_{v'}$ and $V_{v'} \sqsubseteq W_{v'}$. Let $\mathbf{c}_V : \mathcal{C}V_{v'} \rightarrow \mathcal{C}V_v$ and $\mathbf{c}_W : \mathcal{C}W_{v'} \rightarrow \mathcal{C}W_w$ be comparison maps. Then define

$$\rho_{[W]}^{[V]} = \mathbf{c}_W \left(\rho_{W_{v'}}^{V_{v'}} \right),$$

which is a uniformly bounded set, and define $\rho_{[V]}^{[W]} : \mathcal{C}[W] \rightarrow \mathcal{C}[V]$ by

$$\rho_{[V]}^{[W]} = \mathbf{c}_V \circ \rho_{V_{v'}}^{W_{v'}} \circ \bar{\mathbf{c}}_W,$$

where $\bar{\mathbf{c}}_W$ is a quasi-inverse of \mathbf{c}_W and $\rho_{V_{v'}}^{W_{v'}} : \mathcal{C}W_{v'} \rightarrow \mathcal{C}V_{v'}$ is the map provided by Definition 1.1.(2). Similarly, if $[V] \triangleright [W]$, and there exists $w \in T$ so that \mathfrak{S}_w contains representatives V_w, W_w of $[V], [W]$, then let

$$\rho_{[W]}^{[V]} = \mathbf{c}_W \left(\rho_{W_w}^{V_w} \right).$$

Otherwise, choose a closest pair v, w so that \mathfrak{S}_v (respectively \mathfrak{S}_w) contains a representative of $[V]$ (respectively $[W]$). Let e be the first edge of the geodesic in T joining v to w , so $v = e^-$ (say). Let S be the \sqsubseteq -maximal element of \mathfrak{S}_e , and let

$$\rho_{[V]}^{[W]} = \mathbf{c}_V \left(\rho_{V_e^-}^{\phi_{e^-}^\diamond(S)} \right).$$

This is well-defined by hypothesis (4).

For each \sim -class $[W]$, let $\rho_T^{[W]}$ be the support of $[W]$ (a uniformly bounded set since \mathcal{T} has bounded supports). Define $\rho_{[W]}^T : T \rightarrow \mathcal{C}W$ as follows: given $v \in T$ not in the support of $[W]$, let e be the unique edge with e^- (say) separating v from the support of $[W]$. Let $S \in \mathfrak{S}_e$ be \sqsubseteq -maximal. Then $\rho_{[W]}^T(v) = \rho_{[W]}^{\phi_{e^-}^\diamond(S)}$. If v is in the support of $[W]$, then let $\rho_{[W]}^T(v)$ be chosen arbitrarily.

Finally, let $K^\perp, K^{\perp'} \in \bigcup_\eta \mathcal{K}_\eta$ and let $[W]$ be a \sim -class. If $K^\perp \pitchfork K^{\perp'}$, then $\rho_{K^\perp}^{K^{\perp'}}$ is the single point $\mathcal{C}K^\perp$. If $K^\perp \sqsubseteq K^{\perp'}$, then $\rho_{K^\perp}^{K^{\perp'}}$ is a constant map and $\rho_{K^\perp}^{K^{\perp'}}$ is the obvious single point. We never have $K^\perp \sqsubseteq [W]$. If $[W] \sqsubseteq K^\perp$, then $\rho_{K^\perp}^{[W]}$ is again the obvious single point, and we can define $\rho_{[W]}^{K^\perp}: \mathcal{C}K^\perp \rightarrow \mathcal{C}[W]$ to be an arbitrary constant map. Finally, $\rho_{K^\perp}^T: T \rightarrow \mathcal{C}K^\perp$ is the constant map, and $\rho_T^{K^\perp}$ is the bounded set defined as follows. By definition, there is a unique pairwise orthogonal set $[W_1], \dots, [W_\eta]$ so that $K^\perp = K_\eta^\perp([W_1], \dots, [W_\eta])$. By the proof of Lemma 8.13, the supports of the various $[W_i]$ all intersect in a subtree of T , which necessarily has diameter at most n , by the bounded supports hypothesis; we take this intersection to be $\rho_T^{K^\perp}$.

We are now ready to complete the proof of the combination theorem.

Proof of Theorem 8.6. We claim that $(\mathcal{X}(\mathcal{T}), \mathfrak{S})$ is hierarchically hyperbolic. We take the nesting, orthogonality, and transversality relations for a tree of spaces given by Definition 8.7 and Definition 8.11. In Lemmas 8.9, 8.10, and Remark 8.12, it is shown that these relations satisfy all of the conditions (2) and (3) of Definition 1.1 not involving the projections. Moreover, the complexity of $(\mathcal{X}(\mathcal{T}), \mathfrak{S})$ is finite, by Lemma 8.13, verifying Definition 1.1.(5). The set of δ -hyperbolic spaces $\{\mathcal{C}A : A \in \mathfrak{S}\}$ is provided by Definition 8.14, while the projections $\pi_{[W]}: \mathcal{X} \rightarrow \mathcal{C}[W]$ required by Definition 1.1.(1) are defined in Definition 8.19 and are uniformly coarsely Lipschitz by Lemma 8.22. Since $\pi_{[W]}(\mathcal{X})$ uniformly coarsely coincides with the image of an appropriately-chosen vertex space \mathcal{X}_v , it is quasiconvex since $\pi_{W_v}(\mathcal{X}_v)$ is uniformly quasiconvex by Definition 1.1.(1). The projections $\rho_{[W]}^{[V]}$ when $[V], [W]$ are non-orthogonal are described in Definition 8.23. To complete the proof, it thus suffices to verify the consistency inequalities (Definition 1.1.(4)), the bounded geodesic image axiom (Definition 1.1.(7)), the large link axiom (Definition 1.1.(6)), partial realization (Definition 1.1.(8)), and uniqueness (Definition 1.1.(9)).

Consistency: Any consistency inequalities involving elements of $\bigcup_\eta \mathcal{K}_\eta$ hold trivially since in that case, at least one of the two associated hyperbolic spaces in question is a point. Suppose that $[U] \pitchfork [V]$ or $[U] \sqsubseteq [V]$ and let $x \in \mathcal{X}$. Choose representatives $U_u \in \mathfrak{S}_u, V_v \in \mathfrak{S}_v$ of $[U], [V]$ so that $d_T(u, v)$ realizes the distance between the supports of $[U], [V]$. By composing the relevant maps in the remainder of the argument with comparison maps, we can assume that U_u, V_v are favorite representatives. Without loss of generality, there exists a vertex $w \in T$ so that $x \in \mathcal{X}_w$. If $u = v$, then consistency follows since it holds in each vertex space, so assume that u, v have disjoint supports and in particular $[U] \pitchfork [V]$.

If $w \notin [u, v]$, then (say) u separates w from v . Then $\pi_{[V]}(x) = \pi_{V_v}(\mathfrak{g}_v(x))$. Let e be the edge of the geodesic $[u, v]$ emanating from u , so that $\rho_{[U]}^{[V]} = \rho_{U_u}^S$, where S is the image in \mathfrak{S}_u of the \sqsubseteq -maximal element of \mathfrak{S}_e . If $d_{CU_u}(\mathfrak{g}_u(x), \rho_{U_u}^S) \leq E$, then we are done. Otherwise, by consistency in \mathfrak{S}_u , we have $d_{CS}(\mathfrak{g}_u(x), \rho_S^{U_u}) \leq E$, from which consistency follows. Hence suppose that $w \in [u, v]$. Then without loss of generality, there is an edge e containing v and separating w from u . As before, projections to V factor through the \sqsubseteq -maximal element of \mathfrak{S}_e , from which consistency follows.

We verify consistency for $T, [W]$ for each \sim -class $[W]$. Choose $x \in \mathcal{X}_v$. If $d_T(v, T_{[W]}) \geq n + 1$, then let e be the edge incident to $T_{[W]}$ separating it from v , so that (up to a comparison map) $\rho_{[W]}^T(v) = \rho_W^S$, where $W \in \mathfrak{S}_{e^+}$ represents W , and $e^+ \in T_{[W]}$, and S is the image in \mathfrak{S}_{e^+} of the \sqsubseteq -maximal element of \mathfrak{S}_e . (Note that $W \pitchfork S$ by hypothesis (4) of the theorem and the choice of e). On the other hand (up to a comparison map) $\pi_{[W]}(x) = \pi_W(\mathfrak{g}_{e^+}(x)) \asymp \pi_W(\mathfrak{g}_{e^+}(\mathcal{X}_v)) \asymp \pi_W(\mathbf{F}_S) = \rho_W^S$, as desired. (The final coarse equality holds by hypothesis (4).)

Finally, suppose that $[U] \sqsubseteq [V]$ and that either $[V] \sqsubset [W]$ or $[V] \pitchfork [W]$ and $[U] \not\sqsubseteq [W]$. We claim that $d_{[W]}(\rho_{[W]}^{[U]}, \rho_{[W]}^{[V]})$ is uniformly bounded. By definition, $T_{[U]} \cap T_{[V]} \neq \emptyset$ and we fix representatives $U_u \in \mathfrak{S}_u, V_u \in \mathfrak{S}_u$ of $[U], [V]$ with $U_u \sqsubseteq V_u$.

Next, suppose that $T_{[V]} \cap T_{[W]} \neq \emptyset$ and $T_{[U]} \cap T_{[W]} \neq \emptyset$, so that we can choose vertices v, w of T and representatives $V_w, W_w \in \mathfrak{S}_w$ so that $V_w \sqsubseteq W_w$ or $V_w \pitchfork W_w$ according to whether $[V] \sqsubseteq [W]$ or $[V] \pitchfork [W]$, and choose representatives $U_v, W_v \in \mathfrak{S}_v$ of $[U], [W]$ so that $U_v \sqsubseteq W_v$ or $U_v \pitchfork W_v$ according to whether $[U] \sqsubseteq [W]$ or $[U] \pitchfork [W]$. Let $m \in T$ be the median of u, v, w . Since u, w lie in the support of $[U], [W]$, so does m , since supports are connected. Likewise, m lies in the support of $[V]$. Let U_m, V_m, W_m be the representatives of $[U], [V], [W]$ in m . Since edge-maps are full hieromorphisms, we have $U_m \sqsubseteq V_m$ and $U_m \not\sqsubseteq W_m$ and either $V_m \sqsubseteq W_m$ or $V_m \pitchfork W_m$. Hence Definition 1.1.(4) implies that $d_{W_m}(\rho_{W_m}^{U_m}, \rho_{W_m}^{V_m})$ is uniformly bounded. Since the comparison maps are uniform quasi-isometries, it follows that $d_{[W]}(\rho_{[W]}^{[U]}, \rho_{[W]}^{[V]})$ is uniformly bounded, as desired.

Next, suppose that $T_{[U]} \cap T_{[W]} = \emptyset$. Then $[U] \pitchfork [W]$. If there is an edge e separating $T_{[W]}$ from $T_{[U]} \cup T_{[V]}$, then $\rho_{[W]}^{[U]} = \rho_{[W]}^{[V]}$ by definition. Otherwise, $[V] \pitchfork [W]$ (by transitivity of \sqsubseteq and the fact that $T_{[U]} \cap T_{[W]} = \emptyset$) but there is some $v \in T_{[V]} \cap T_{[W]}$ and representatives $V_v, W_v \in \mathfrak{S}_v$ of $[V], [W]$ with $V_v \pitchfork W_v$. But by fullness of the hieromorphism and induction on $d_T(u, v)$, we find that $v \in T_{[U]}$, contradicting that $T_{[U]} \cap T_{[W]} = \emptyset$.

Bounded geodesic image and large link axiom in T and $\bigcup_{\eta} \mathcal{K}_{\eta}$: For bounded geodesic image, in the case where one of the two nested elements of \mathfrak{S} in question is in $\bigcup_{\eta} \mathcal{K}_{\eta}$, the claim holds trivially since the associated hyperbolic space has diameter 0 and the associated map between hyperbolic spaces has either domain or codomain a single point.

Let γ be a geodesic in T and let $[W]$ be a \sim -class so that $d_T(\gamma, \rho_T^{[W]}) > 1$, which is to say that γ does not contain vertices in the support of $[W]$. Let e be the terminal edge of the geodesic joining γ to the support of $[W]$. Then for all $u \in \gamma$, we have by definition that $\rho_{[W]}^T(u) = \rho_{[W]}^{[S]}$ for some fixed \sim -class $[S]$. This verifies the bounded geodesic image axiom for $T, [W]$.

By Lemma 8.17, there exists a constant K'' such that if $x, x' \in \mathcal{X}$ respectively project to vertices v, v' , then any $[W] \in \mathfrak{S} - \{T\}$ with $d_{[W]}(x, x') \geq K''$ is supported on a vertex $v_{[W]}$ on the geodesic $[v, v']$ and is hence nested into $[S_{v_w}]$, where S_{v_w} is maximal in $\mathfrak{S}_{v_{[W]}}$. Indeed, choose w in the support of $[W]$. Then either $d_{[W]}(x, x')$ is smaller than some specified constant, or $d_{\mathcal{X}_w}(\mathfrak{g}_w(x), \mathfrak{g}_w(x')) > K'$. Thus $\mathfrak{g}_{\mathcal{X}_w}(\mathcal{X}_m)$ has diameter at least K' , where m is the median of v, v', w . Hence m lies in the support of $[W]$, and $m \in [v, v']$, and $[W] \sqsubseteq [S]$, where S is \sqsubseteq -maximal in \mathfrak{S}_m . Finally, for each such S_{v_w} , it is clear that $d_T(x, \rho_T^{[S_{v_w}]}) \leq d_T(x, x')$, verifying the conclusion of the large link axiom for T .

Finally, we have to check that if $K^{\perp} = K^{\perp}([W_1], \dots, [W_n])$ and $x, y \in \mathcal{X}$, then there exist a uniformly bounded number of elements $[U_j]$ so that for any $[V] \sqsubseteq K^{\perp}$ with $d_{[V]}(x, y) \geq E$, the class $[V]$ is nested into some U_j . We have shown above that any such $[V]$ is nested into the \sqsubseteq -maximal $S_v \in \mathfrak{S}_v$ for some v on the geodesic of T between $\pi_T(x)$ and $\pi_T(y)$. Now, since $[V] \sqsubseteq K^{\perp}$, we have $[V] \perp [W_i]$ for all i , so that the support of $[V]$ uniformly coarsely coincides with the support of $[W_i]$ for each i . Hence v must be among the boundedly many vertices on $[\pi_T(x), \pi_T(y)]$ that lie in the intersection of the supports of the $[W_i]$. Thus we can take our set of U_j to be the set of such S_v , which has uniformly bounded cardinality (bounded in terms of n).

Bounded geodesic image and large link axiom in $W \sqsubset T$: Let $[W]$ be non- \sqsubseteq -maximal, let $[V] \sqsubseteq [W]$, and let γ be a geodesic in $\mathcal{C}[W]$. Then γ is a geodesic in $\mathcal{C}W_w$, by definition, where w is the favorite vertex of $[W]$ with corresponding representative W_w . Let

V_w be the representative of $[V]$ supported on w , so that $\rho_{[W]}^{[V]} = \rho_{W_w}^{V_w}$, so that γ avoids the E -neighborhood of $\rho_{[W]}^{[V]}$ exactly when it avoids the E -neighborhood of $\rho_{W_w}^{V_w}$. The bounded geodesic image axiom now follows from bounded geodesic image in \mathfrak{S}_w , although the constant E has been changed according to the quasi-isometry constant of comparison maps.

Now suppose $x, x' \in \mathcal{X}_v, \mathcal{X}_{v'}$ and choose w to be the favorite vertex in the support of $[W]$. Suppose for some $[V] \subseteq [W]$ that $d_{[V]}(x, x') \geq E'$, where E' depends on E and the quasi-isometry constants of the edge-hieromorphisms. Then $d_{V_w}(\mathfrak{g}_w(x), \mathfrak{g}_w(x')) \geq E$, for some representative $V_w \in \mathfrak{S}_w$ of $[V]$, by our choice of E' . Hence, by the large link axiom in \mathfrak{S}_w , we have that $V_w \subseteq T_i$, where $\{T_i\}$ is a specified set of $N = [d_{W_w}(\mathfrak{g}_w(x), \mathfrak{g}_w(x'))] = [d_{[W]}(x, x')]$ elements of \mathfrak{S}_w , with each $T_i \subset W_w$. Moreover, the large link axiom in \mathfrak{S}_w implies that $d_{[W]}(x, \rho_{[W]}^{[T_i]}) = d_{W_w}(\mathfrak{g}_w(x), \rho_{W_w}^{T_i}) \leq N$ for all i . This verifies the large link axiom for $(\mathcal{X}(\mathcal{T}), \mathfrak{S})$.

Partial realization: Let $[V_1], \dots, [V_k] \in \mathfrak{S}$ be pairwise-orthogonal, and, for each $i \leq k$, let $p_i \in \mathcal{C}[V_i]$. For each i , let $T_i \subseteq T$ be the induced subgraph spanned by the vertices w such that $[V_i]$ has a representative in \mathfrak{S}_w . The definition of the \sim -relation implies that each T_i is connected, so by the Helly property of trees, there exists a vertex $v \in T$ such that for each i , there exists $V_v^i \in \mathfrak{S}_v$ representing $[V_i]$. Moreover, we have $V_v^i \perp V_v^j$ for $i \neq j$, since the edge-hieromorphisms preserve the orthogonality relation. Applying the partial realization axiom (Definition 1.1.(8)) to $\{p'_i \in \mathcal{C}V_v^i\}$, where p'_i is the image of p_i under the appropriate comparison map, yields a point $x \in \mathcal{X}_v$ such that $\pi_{V_v^i}(x)$ is coarsely equal to p'_i for all i , whence $d_{[V_i]}(x, p_i)$ is uniformly bounded. If $[V_i] \subseteq [W]$, then W has a representative $V_w \in \mathfrak{S}_v$ such that $V_w^i \subseteq W$, whence $d_{[W]}(x, \rho_{[W]}^{[V_i]})$ is uniformly bounded since x is a partial realization point for $\{V_w^i\}$ in \mathfrak{S}_v . Finally, if $[W] \not\subseteq [V_i]$, then either the subtrees of T supporting $[W]$ and $[V_i]$ are disjoint, in which case $d_{[W]}(x, \rho_{[W]}^{[V_i]})$ is bounded, or $[W]$ has a representative in \mathfrak{S}_v transverse to V_v^i , in which case the same inequality holds by our choice of x . There is nothing to check regarding projections onto $\mathcal{C}K^\perp$ for $K^\perp \in \bigcup_\eta \mathcal{K}_\eta$, since those spaces are single points.

It remains to consider pairwise orthogonal collections that include elements of $\bigcup_\eta \mathcal{K}_\eta$. Since no two of these elements can be orthogonal, we must consider $K^\perp = K^\perp_\eta([W_1], \dots, [W_\eta])$, which is orthogonal to \sim -classes $[V_1], \dots, [V_k]$, which themselves form an orthogonal collection. Let p be the unique point of $\mathcal{C}K^\perp$, and for each $i \leq k$, let $p_i \in \mathcal{C}V_i$. Then the previous discussion provides a point x so that for any i , we have $\pi_{[V_i]}(x) \asymp p_i$, and $\pi_T(x) \asymp \rho_T^{[V_i]}$. Moreover, for any $[U]$ so that, for some i , we have $[U] \not\subseteq [V_i]$ or $[V_i] \subseteq [U]$, we have $\pi_{[U]}(x) \asymp \rho_{[U]}^{[V_i]}$. We claim that x also satisfies the conclusion of partial realisation for the pairwise-orthogonal set $K^\perp, [V_1], \dots, [V_k]$. Again, there is nothing to check regarding projections onto $\mathcal{C}K^\perp$ for $K^\perp \in \bigcup_\eta \mathcal{K}_\eta$, since those spaces are single points, and this includes the statement about $\pi_{K^\perp}(x)$. So, it just remains to check that $\pi_T(x)$ uniformly coarsely coincides with $\rho_T^{K^\perp}$. But $\pi_T(x)$ coarsely coincides with $\rho_T^{[V_i]}$ for any i , by the construction of x . Since $[V_i] \perp K^\perp$, we have $[V_i] \subseteq [W_j]$ for some j , so $\rho_T^{[V_i]}$ coarsely coincides with $\rho_T^{[W_j]}$. But $\rho_T^{[W_j]}$ coarsely coincides, by definition, with $\rho_T^{K^\perp}$, as required.

Uniqueness of realization points: Suppose $x, y \in \mathcal{X}$ satisfy $d_{[V]}(x, y) \leq K$ for all $[V] \in \mathfrak{S}$. Then, for each vertex $v \in T$, applying the uniqueness axiom in \mathcal{X}_v to $\mathfrak{g}_v(x), \mathfrak{g}_v(y)$ shows that $d_{\mathcal{X}_v}(\mathfrak{g}_v(x), \mathfrak{g}_v(y)) \leq \zeta = \zeta(K)$. Indeed, otherwise we would have $d_V(\mathfrak{g}_v(x), \mathfrak{g}_v(y)) > \xi K + \xi$ for some $V \in \mathfrak{S}_v$, whence $d_{[V]}(x, y) > K$. There exists $k \leq K$ and a sequence v_0, \dots, v_k of vertices in T so that $x \in \mathcal{X}_{v_0}, y \in \mathcal{X}_{v_k}$. For each j , let $x_j = \mathfrak{g}_{v_j}(x)$ and $y_j = \mathfrak{g}_{v_j}(y)$. Then $x = x_0, y_0, x_1, y_1, \dots, y_{j-1}, x_j, y_j, \dots, x_k, y_k = y$ is a path of uniformly

bounded length joining x to y . Indeed, $d_{\mathcal{X}_{v_j}}(x_j, y_j) \leq \zeta$ and $k \leq K$ by the preceding discussion, while x_j coarsely coincides with a point on the opposite side of an edge-space from y_{j-1} by the definition of the gate of an edge-space in a vertex-space and the fact that x_{j-1} and y_{j-1} coarsely coincide. This completes the proof. \square

8.1. Equivariant definition of $(\mathcal{X}(\mathcal{T}), \mathfrak{S})$. Let \mathcal{T} denote the tree of hierarchically hyperbolic spaces $(\mathcal{T}, \{\mathcal{X}_v\}, \{\mathcal{X}_e\}, \{\pi_{e^\pm}\})$, and let $(\mathcal{X}(\mathcal{T}), \mathfrak{S})$ be the hierarchically hyperbolic structure defined in the proof of Theorem 8.6. Various arbitrary choices were made in defining the constituent hyperbolic spaces and projections in this hierarchically hyperbolic structure, and we now insist on a specific way of making these choices in order to describe automorphisms of $(\mathcal{X}(\mathcal{T}), \mathfrak{S})$.

Recall that an automorphism of $(\mathcal{X}(\mathcal{T}), \mathfrak{S})$ is determined by a bijection $g: \mathfrak{S} \rightarrow \mathfrak{S}$ and a set of isometries $g: \mathcal{C}Q \rightarrow \mathcal{C}gQ$, for $Q \in \mathfrak{S}$. Via the distance formula, this determines a uniform quasi-isometry $\mathcal{X}(\mathcal{T}) \rightarrow \mathcal{X}(\mathcal{T})$.

A bijection $g: \bigsqcup_{v \in \mathcal{V}} \mathfrak{S}_v \rightarrow \bigsqcup_{v \in \mathcal{V}} \mathfrak{S}_v$ is T -coherent if there is an induced isometry g of the underlying tree, T , so that $fg = gf$, where $f: \bigsqcup_{v \in \mathcal{V}} \mathfrak{S}_v \rightarrow T$ sends each $V \in \mathfrak{S}_v$ to v , for all $v \in \mathcal{V}$. The T -coherent bijection g is said to be \mathcal{T} -coherent if it also preserves the relation \sim . Noting that the composition of \mathcal{T} -coherent bijections is \mathcal{T} -coherent, denote by $\mathcal{P}_{\mathcal{T}}$ the group of \mathcal{T} -coherent bijections. For each $g \in \mathcal{P}_{\mathcal{T}}$, there is an induced bijection $g: \mathfrak{S}^0 \rightarrow \mathfrak{S}^0$.

Recall that the hierarchically hyperbolic structure $(\mathcal{X}(\mathcal{T}), \mathfrak{S})$ was completely determined except for the following three types of choices which were made arbitrarily.

- (1) For each $[V] \in \mathfrak{S}$, we chose an arbitrary *favorite vertex* v in the support of $[V]$; and,
- (2) we chose an arbitrary *favorite representative* $V_v \in \mathfrak{S}_v$ with $[V] = [V_v]$. (Note that if, as is often the case in practice, edge-hieromorphisms $\mathfrak{S}_e \rightarrow \mathfrak{S}_v$ are injective, then V_v is the unique representative of its \sim -class that lies in \mathfrak{S}_v , and hence our choice is uniquely determined.)
- (3) For each $[W] \in \mathfrak{S}$, the point $\rho_{[W]}^T(v)$ is chosen arbitrarily in $\mathcal{C}W$, where W is the favorite representative of $[W]$ and v is a vertex in the support of $[W]$.

We now constrain these choices so that they are equivariant. For each $\mathcal{P}_{\mathcal{T}}$ -orbit in \mathfrak{S} , choose a representative $[V]$ of that orbit, choose a favorite vertex v arbitrarily in its support, and choose a favorite representative $V_v \in \mathfrak{S}_v$ of $[V]$. Then declare $gV_v \in \mathfrak{S}_{gv}$ to be the favorite representative, and gv the favorite vertex, associated to $g[V]$, for each $g \in \mathcal{P}_{\mathcal{T}}$.

Recall that, for each $[W] \in \mathfrak{S}$, we defined $\mathcal{C}[W]$ to be $\mathcal{C}W$, where W is the favorite representative of $[W]$. Suppose that we have specified a subgroup $G \leq \mathcal{P}_{\mathcal{T}}$ and, for each $[W] \in \mathfrak{S}$ and $g \in \mathcal{P}_{\mathcal{T}}$, an isometry $g: \mathcal{C}[W] \rightarrow \mathcal{C}g[W]$. Then we choose $\rho_{[W]}^T$ in such a way that $\rho_{g[W]}^T = g\rho_{[W]}^T$ for each $[W] \in \mathfrak{S}$ and $g \in G$.

8.2. Graphs of hierarchically hyperbolic groups. Recall that the finitely generated group G is *hierarchically hyperbolic* if there is a hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$ such that $G \leq \text{Aut}(\mathfrak{S})$ and the action of G on \mathcal{X} is metrically proper and cobounded and the action of G on \mathfrak{S} is co-finite (this, together with the definition of an automorphism, implies that only finitely many isometry types of hyperbolic space are involved in the HHS structure). Endowing G with a word-metric, we see that (G, \mathfrak{S}) is a hierarchically hyperbolic space.

If (G, \mathfrak{S}) and (G', \mathfrak{S}') are hierarchically hyperbolic groups, then a *homomorphism of hierarchically hyperbolic groups* $\phi: (G, \mathfrak{S}) \rightarrow (G', \mathfrak{S}')$ is a homomorphism $\phi: G \rightarrow G'$ that is also a ϕ -equivariant hieromorphism as in Definition 1.22.

Recall that a *graph of groups* \mathcal{G} is a graph $\Gamma = (V, E)$ together with a set $\{G_v : v \in V\}$ of vertex groups, a set $\{G_e : e \in E\}$ of edge groups, and monomorphisms $\phi_e^\pm: G_e \rightarrow G_{e^\pm}$, where e^\pm are the vertices incident to e . As usual, the *total group* G of \mathcal{G} is the quotient

of $(*_{v \in V} G_v) * F_E$, where F_E is the free group generated by E , obtained by imposing the following relations:

- $e = 1$ for all $e \in E$ belonging to some fixed spanning tree T of Γ ;
- $\phi_e^+(g) = e\phi_e^-(g)e^{-1}$ for $e \in E$ and $g \in G_e$.

We are interested in the case where Γ is a finite graph and, for each $v \in V, e \in E$, we have sets $\mathfrak{S}_v, \mathfrak{S}_e$ so that (G_v, \mathfrak{S}_v) and (G_e, \mathfrak{S}_e) are hierarchically hyperbolic group structures for which $\phi_e^\pm: G_e \rightarrow G_{e^\pm}$ is a homomorphism of hierarchically hyperbolic groups. In this case, \mathcal{G} is a *finite graph of hierarchically hyperbolic groups*. If in addition each ϕ_e^\pm has hierarchically quasiconvex image, then \mathcal{G} has *quasiconvex edge groups*.

Letting $\tilde{\Gamma}$ denote the Bass-Serre tree, observe that $\mathcal{T} = \tilde{\mathcal{G}} = (\tilde{\Gamma}, \{G_{\tilde{v}}\}, \{G_{\tilde{e}}\}, \{\phi_{\tilde{e}}^\pm\})$ is a tree of hierarchically hyperbolic spaces, where \tilde{v} ranges over the vertex set of $\tilde{\Gamma}$, and each $G_{\tilde{v}}$ is a conjugate in the total group G to G_v , where $\tilde{v} \mapsto v$ under $\tilde{\Gamma} \rightarrow \Gamma$, and an analogous statement holds for edge-groups. Each $\phi_{\tilde{e}}^\pm$ is conjugate to an edge-map in \mathcal{G} in the obvious way. We say \mathcal{G} has *bounded supports* if \mathcal{T} does.

Corollary 8.24 (Combination theorem for HHGs). *Let $\mathcal{G} = (\Gamma, \{G_v\}, \{G_e\}, \{\phi_e^\pm\})$ be a finite graph of hierarchically hyperbolic groups, with Bass-Serre tree $\tilde{\Gamma}$. Suppose that:*

- (1) \mathcal{G} has quasiconvex edge groups;
- (2) each ϕ_e^\pm , as a hieromorphism, is full;
- (3) \mathcal{G} has bounded supports;
- (4) if e is an edge of Γ and S_e the \sqsubseteq -maximal element of \mathfrak{S}_e , then for all $V \in \mathfrak{S}_{e^\pm}$, the elements V and $\phi_{e^\pm}^\diamond(S_e)$ are not orthogonal in \mathfrak{S}_{e^\pm} ;
- (5) for each vertex v of Γ , there are finitely many G_v -orbits of subsets $\mathcal{U} \subset \mathfrak{S}_v$ for which the elements of \mathcal{U} are pairwise-orthogonal;
- (6) there exists $K \geq 0$ such that for all vertices v of $\tilde{\Gamma}$ and edges e incident to v , we have $d_{\text{Haus}}(\phi_v(G_e), \mathbf{F}_{\phi_v^\diamond(S_e)} \times \{\star\}) \leq K$, where $S_e \in \mathfrak{S}_e$ is the unique maximal element and $\star \in \mathbf{E}_{\phi_v^\diamond(S_e)}$.

Then the total group G of \mathcal{G} is a hierarchically hyperbolic group.

Remark 8.25. We have added hypothesis (5) because it is exactly what's required. In fact, it should follow from a stronger but more natural condition, namely that for each $V \in \mathfrak{S}_v$, the stabiliser in G_v of the standard product region \mathbf{P}_V acts cocompactly on \mathbf{P}_V . This holds, for example, in the mapping class group. On the other hand, this stronger condition it is not a consequence of the definition of an HHG since, for example, one can put exotic HHG structures on a free group where this fails.

Proof of Corollary 8.24. By Theorem 8.6, (G, \mathfrak{S}) is a hierarchically hyperbolic space. Observe that $G \leq \mathcal{P}_{\mathcal{G}}$, since G acts on the Bass-Serre tree $\tilde{\Gamma}$, and this action is induced by an action on $\bigcup_{v \in V} \mathfrak{S}_v$ preserving the \sim -relation. Hence the hierarchically hyperbolic structure (G, \mathfrak{S}) can be chosen according to the constraints in Section 8.1, whence it is easily checked that G acts on \mathfrak{S}^0 by HHS automorphisms. Moreover, for any $[V]$, there are finitely many $\text{Stab}_G([V])$ -orbits of \sim -classes nested in $[V]$.

The action on \mathfrak{S}^0 is co-finite since each G_v is a hierarchically hyperbolic group.

Moreover, since G preserves nesting and orthogonality in \mathfrak{S}^0 , we have an induced action of G on $\bigcup_\eta \mathcal{K}_\eta$ defined by $gK^\perp([W_1], \dots, [W_\eta]) = K^\perp([gW_1], \dots, [gW_\eta])$. We must show that this action (and hence the action of G on \mathfrak{S} obtained by combining this with the action on \mathfrak{S}^0) is cofinite.

Since each element of \mathcal{K}_n corresponds to a n -element pairwise-orthogonal set in \mathfrak{S}^0 , and this correspondence is injective, it suffices to show that there are only finitely many G -orbits of such sets. This follows from hypothesis (5).

Finally, the maps of the form $\pi_{K^\perp} : G \rightarrow CK^\perp$ and $\rho_{K^\perp}^*$ obviously satisfy the conditions required of an action by HHS automorphisms, since they are constant maps. Finally, for all $K^\perp \in \bigcup_\eta \mathcal{K}_\eta$ and $[W]$, and $g \in G$, we can choose the arbitrary constant map $\rho_{[W]}^{K^\perp}$ so that $g(\rho_{[W]}^{K^\perp}) = \rho_{g[W]}^{gK^\perp}$, where $g : \mathcal{C}[W] \rightarrow \mathcal{C}g[W]$ is the isometry from the automorphism action on \mathfrak{S}^0 . The same holds with W replaced by T , since $\rho_T^{K^\perp}$ was defined to be an intersection of support trees associated to K^\perp , and $\rho_T^{gK^\perp}$ is, by the definition of the G -action on $\bigcup_\eta \mathcal{K}_\eta$ and the G -equivariance of the assignment of each \sim -class to its support tree, the intersection of the g -translates of these support trees, i.e. $g\rho_T^{K^\perp}$. This completes the proof. \square

Remark 8.26 (Examples where the combination theorem does not apply). Examples where one cannot apply Theorem 8.6 or Corollary 8.24 are likely to yield examples of groups that are not hierarchically hyperbolic groups, or even hierarchically hyperbolic spaces.

- (1) Let \mathcal{G} be a finite graph of groups with \mathbb{Z}^2 vertex groups and \mathbb{Z} edge groups, i.e., a *tubular group*. In [Wis14], Wise completely characterized the tubular groups that act freely on CAT(0) cube complexes, and also characterized the (rare) tubular groups that admit cocompact such actions; Woodhouse recently gave a necessary and sufficient condition for the particular cube complex constructed in [Wis14] to be finite-dimensional [Woo16]. These results suggest that there is little hope of producing hierarchically hyperbolic structures for tubular groups via cubulation, except in particularly simple cases.

This is because the obstruction to cocompact cubulation is very similar to the obstruction to building a hierarchically hyperbolic structure using Theorem 8.6. Indeed, if some vertex-group $G_v \cong \mathbb{Z}^2$ has more than 2 independent incident edge-groups, then, if \mathcal{G} satisfied the hypotheses of Theorem 8.6, the hierarchically hyperbolic structure on G_v would include 3 pairwise-orthogonal unbounded elements, contradicting partial realization. This shows that such a tubular group does not admit a hierarchically hyperbolic structure by virtue of the obvious splitting, and in fact shows that there is no hierarchically hyperbolic structure in which G_v and the incident edge-groups are hierarchically quasiconvex.

- (2) Let $G = F \rtimes_\phi \mathbb{Z}$, where F is a finite-rank free group and $\phi : F \rightarrow F$ an automorphism. When F is atoroidal, G is a hierarchically hyperbolic group simply by virtue of being hyperbolic [BF92, Bri00]. There is also a more refined hierarchically hyperbolic structure in this case, in which all of the hyperbolic spaces involved are quasi-trees. Indeed, by combining results in [HW15] and [Ago13], one finds that G acts freely, cocompactly, and hence virtually co-specially on a CAT(0) cube complex, which therefore contains a G -invariant *factor system* in the sense of [BHS17a] and is hence a hierarchically hyperbolic group; the construction in [BHS17a] ensures that the hierarchically hyperbolic structure for such cube complexes always uses a collection of hyperbolic spaces uniformly quasi-isometric to trees. However, the situation is presumably quite different when G is not hyperbolic. In this case, it seems that G is rarely hierarchically hyperbolic.

8.3. Products. In this short section, we briefly describe a hierarchically hyperbolic structure on products of hierarchically hyperbolic spaces.

Proposition 8.27 (Product HHS). *Let $(\mathcal{X}_0, \mathfrak{S}_0)$ and $(\mathcal{X}_1, \mathfrak{S}_1)$ be hierarchically hyperbolic spaces. Then $\mathcal{X} = \mathcal{X}_0 \times \mathcal{X}_1$ admits a hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{S})$ such that for each of $i \in \{0, 1\}$ the inclusion map $\mathcal{X}_i \rightarrow \mathcal{X}$ induces a quasiconvex hieromorphism.*

Proof. Let $(\mathcal{X}_i, \mathfrak{S}_i)$ be hierarchically hyperbolic spaces for $i \in \{0, 1\}$. Let \mathfrak{S} be a hierarchically hyperbolic structure consisting of the disjoint union of \mathfrak{S}_0 and \mathfrak{S}_1 (together with

their intrinsic hyperbolic spaces, projections, and nesting, orthogonality, and transversality relations), along with the following domains whose associated hyperbolic spaces are points: S , into which everything will be nested; U_i , for $i \in \{0, 1\}$, into which everything in \mathfrak{S}_i is nested; for each $U \in \mathfrak{S}_i$ a domain V_U , with $|\mathcal{C}V_U| = 1$, into which is nested everything in \mathfrak{S}_{i+1} and everything in \mathfrak{S}_i orthogonal to U . The elements V_U are all transverse to U_0 and U_1 . Given U, U' , the elements $V_U, V_{U'}$ are transverse unless $U \sqsubseteq U'$, in which case $V_U \sqsubseteq V_{U'}$. Projections $\pi_U: \mathcal{X}_0 \times \mathcal{X}_1 \rightarrow U \in \mathfrak{S}$ are defined in the obvious way when $U \notin \mathfrak{S}_0 \cup \mathfrak{S}_1$; otherwise, they are the compositions of the existing projections with projection to the relevant factor. Projections of the form ρ_V^U are either defined already, uniquely determined, or are chosen to coincide with the projection of some fixed basepoint (when $V \in \mathfrak{S}_0 \cup \mathfrak{S}_1$ and U is not). It is easy to check that this gives a hierarchically hyperbolic structure on $\mathcal{X}_1 \times \mathcal{X}_2$.

The hieromorphisms $(\mathcal{X}_i, \mathfrak{S}_i) \rightarrow (\mathcal{X}, \mathfrak{S})$ are inclusions on \mathcal{X}_i and \mathfrak{S} ; for each $U \in \mathfrak{S}_i$, the map $\mathfrak{S}_i \ni \mathcal{C}U \rightarrow \mathcal{C}U \in \mathfrak{S}$ is the identity. It follows immediately from the definitions that the diagrams from Definition 1.20 coarsely commute, so that these maps are indeed hieromorphisms. Hierarchical quasiconvexity likewise follows from the definition. \square

Product HHS will be used in defining hierarchically hyperbolic structures on graph manifolds in Section 10. The next result follows directly from the proof of the previous proposition.

Corollary 8.28. *Let G_0 and G_1 be hierarchically hyperbolic groups. Then $G_0 \times G_1$ is a hierarchically hyperbolic group.*

9. HYPERBOLICITY RELATIVE TO HHGS

Relatively hyperbolic groups possess natural hierarchically hyperbolic structures:

Theorem 9.1 (Hyperbolicity relative to HHGs). *Let the group G be hyperbolic relative to a finite collection \mathcal{P} of peripheral subgroups. If each $P \in \mathcal{P}$ is a hierarchically hyperbolic space, then G is a hierarchically hyperbolic space. Further, if each $P \in \mathcal{P}$ is a hierarchically hyperbolic group, then so is G .*

Proof. We prove the statement about hierarchically hyperbolic groups; the statement about spaces follows *a fortiori*.

For each $P \in \mathcal{P}$, let (P, \mathfrak{S}_P) be a hierarchically hyperbolic group structure. For convenience, assume that the $P \in \mathcal{P}$ are pairwise non-conjugate (this will avoid conflicting hierarchically hyperbolic structures). For each P and each left coset gP , let \mathfrak{S}_{gP} be a copy of \mathfrak{S}_P (with associated hyperbolic spaces and projections), so that there is a hieromorphism $(P, \mathfrak{S}_P) \rightarrow (gP, \mathfrak{S}_{gP})$, equivariant with respect to the conjugation isomorphism $P \rightarrow P^g$.

Let \widehat{G} be the usual hyperbolic space formed from G by coning off each left coset of each $P \in \mathcal{P}$. Let $\mathfrak{S} = \{\widehat{G}\} \cup \bigsqcup_{gP \in G\mathcal{P}} \mathfrak{S}_{gP}$. The nesting, orthogonality, and transversality relations on each \mathfrak{S}_{gP} are as defined above; if $U, V \in \mathfrak{S}_{gP}, \mathfrak{S}_{g'P'}$ and $gP \neq g'P'$, then declare $U \pitchfork V$. Finally, for all $U \in \mathfrak{S}$, let $U \sqsubseteq \widehat{G}$. The hyperbolic space $\mathcal{C}\widehat{G}$ is \widehat{G} , while the hyperbolic space $\mathcal{C}U$ associated to each $U \in \mathfrak{S}_{gP}$ was defined above.

The projections are defined as follows: $\pi_{\widehat{G}}: G \rightarrow \widehat{G}$ is the inclusion, which is coarsely surjective and hence has quasiconvex image. For each $U \in \mathfrak{S}_{gP}$, let $\mathfrak{g}_{gP}: G \rightarrow gP$ be the closest-point projection onto gP and let $\pi_U = \pi_U \circ \mathfrak{g}_{gP}$, to extend the domain of π_U from gP to G . Since each π_U was coarsely Lipschitz on U with quasiconvex image, and the closest-point projection is uniformly coarsely Lipschitz, the projection π_U is uniformly coarsely Lipschitz and has quasiconvex image. For each $U, V \in \mathfrak{S}_{gP}$, the coarse maps ρ_U^V and ρ_V^U were already defined. If $U \in \mathfrak{S}_{gP}$ and $V \in \mathfrak{S}_{g'P'}$, then $\rho_V^U = \pi_V(\mathfrak{g}_{g'P'}(gP))$, which is a uniformly bounded set (here we use relative hyperbolicity, not just the weak relative hyperbolicity that is all we needed so far). Finally, for $U \neq \widehat{G}$, we define $\rho_{\widehat{G}}^U$ to be the cone-point over the unique

gP with $U \in \mathfrak{S}_{gP}$, and $\rho_U^{\hat{G}}: \hat{G} \rightarrow \mathcal{CU}$ is defined as follows: for $x \in G$, let $\rho_U^{\hat{G}}(x) = \pi_U(x)$. If $x \in \hat{G}$ is a cone-point over $g'P' \neq gP$, let $\rho_U^{\hat{G}}(x) = \rho_U^{S_{g'P'}}$, where $S_{g'P'} \in \mathfrak{S}_{g'P'}$ is \sqsubseteq -maximal. The cone-point over gP may be sent anywhere in U .

By construction, to verify that (G, \mathfrak{S}) is a hierarchically hyperbolic group structure, it suffices to verify that it satisfies the remaining axioms for a hierarchically hyperbolic space given in Definition 1.1, since the additional G -equivariance conditions hold by construction. Specifically, it remains to verify consistency, bounded geodesic image and large links, partial realization, and uniqueness.

Consistency: The nested consistency inequality holds automatically within each \mathfrak{S}_{gP} , so it remains to verify it only for $U \in \mathfrak{S}_{gP}$ versus \hat{G} , but this follows directly from the definition: if $x \in G$ is far in \hat{G} from the cone-point over gP , then $\rho_U^{\hat{G}}(x) = \pi_U(x)$, by definition. To verify the transverse inequality, it suffices to consider $U \in \mathfrak{S}_{gP}, V \in \mathfrak{S}_{g'P'}$ with $gP \neq g'P'$. Let $x \in G$ and let $z = \mathfrak{g}_{g'P'}(x)$. Then, if $d_U(x, z)$ is sufficiently large, then $d_{gP}(x, z)$ is correspondingly large, so that by Lemma 1.15 of [Sis13], $\mathfrak{g}_{g'P'}(x)$ and $\mathfrak{g}_{g'P'}(gP)$ coarsely coincide, as desired.

The last part of the consistency axiom, Definition 1.1.(4), holds as follows. Indeed, if $U \sqsubseteq V$, then either $U = V$, and there is nothing to prove. Otherwise, if $U \sqsubseteq V$ and either $V \not\sqsubseteq W$ or $W \not\sqsupseteq V$, then either $U, V \in \mathfrak{S}_{gP}$ for some g, P , or $U \in \mathfrak{S}_{gP}$ and $V = \hat{G}$. The latter situation precludes the existence of W , so we must be in the former situation. If $W \in \mathfrak{S}_{gP}$, we are done since the axiom holds in \mathfrak{S}_{gP} . If $W = \hat{G}$, then U, V both project to the cone-point over gP , so $\rho_W^U = \rho_W^V$. In the remaining case, $W \in \mathfrak{S}_{g'P'}$ for some $g'P' \neq gP$, in which case ρ_W^U, ρ_W^V both coincide with $\pi_W(\mathfrak{g}_{g'P'}(gP))$.

Bounded geodesic image: Bounded geodesic image holds within each \mathfrak{S}_{gP} by construction, so it suffices to consider the case of $U \in \mathfrak{S}_{gP}$ nested into \hat{G} . Let $\hat{\gamma}$ be a geodesic in \hat{G} avoiding gP and the cone on gP . Lemma 1.15 of [Sis13] ensures that any lift of $\hat{\gamma}$ has uniformly bounded projection on gP , so $\rho_U^{\hat{G}} \circ \hat{\gamma}$ is uniformly bounded.

Large links: The large link axiom (Definition 1.1.(6)) can be seen to hold in (G, \mathfrak{S}) by combining the large link axiom in each gP with malnormality of \mathcal{P} and Lemma 1.15 of [Sis13].

Partial realization: This follows immediately from partial realization within each \mathfrak{S}_{gP} and the fact that no new orthogonality was introduced in defining (G, \mathfrak{S}) , together with the definition of \hat{G} and the definition of projection between elements of \mathfrak{S}_{gP} and $\mathfrak{S}_{g'P'}$ when $gP \neq g'P'$. More precisely, if $U \in \mathfrak{S}_{gP}$ and $p \in \mathcal{CU}$, then by partial realization within gP , there exists $x \in gP$ so that $d_U(x, p) \leq \alpha$ for some fixed constant α and $d_V(x, \rho_V^U) \leq \alpha$ for all $V \in \mathfrak{S}_{gP}$ with $U \sqsubseteq V$ or $U \not\sqsupseteq V$. Observe that $d_{\hat{G}}(x, \rho_{\hat{G}}^U) = 1$, since $x \in gP$ and $\rho_{\hat{G}}^U$ is the cone-point over gP . Finally, if $g'P' \neq gP$ and $V \in \mathfrak{S}_{g'P'}$, then $d_V(x, \rho_V^U) = d_V(\pi_V(\mathfrak{g}_{g'P'}(x)), \pi_V(\mathfrak{g}_{g'P'}(gP))) = 0$ since $x \in gP$.

Uniqueness: If x, y are uniformly close in \hat{G} , then either they are uniformly close in G , or they are uniformly close to a common cone-point, over some gP , whence the claim follows from the uniqueness axiom in \mathfrak{S}_{gP} . \square

Remark 9.2. The third author established a characterization of relative hyperbolicity in terms of projections in [Sis13]. Further, there it was proven that like mapping class groups, there was a natural way to compute distances in relatively hyperbolic groups from certain related spaces, namely: if (G, \mathcal{P}) is relatively hyperbolic, then distances in G are coarsely obtained by summing the corresponding distance in the coned-off Cayley graph \hat{G} together with the distances between projections in the various $P \in \mathcal{P}$ and their cosets. We recover a new proof of Sisto's formula as a consequence of Theorem 9.1 and Theorem 4.5.

Theorem 9.1 will be used in our analysis of 3-manifold groups in Section 10. However, there is a more general statement in the context of metrically relatively hyperbolic spaces (e.g., what Drutu–Sapir call asymptotically tree-graded [DS05], or spaces that satisfy the equivalent condition on projections formulated in [Sis12]). For instance, arguing exactly as in the proof of Theorem 9.1 shows that if the space \mathcal{X} is hyperbolic relative to a collection of uniformly hierarchically hyperbolic spaces, then \mathcal{X} admits a hierarchically hyperbolic structure (in which each peripheral subspace embeds hieromorphically).

More generally, let the geodesic metric space \mathcal{X} be hyperbolic relative to a collection \mathcal{P} of subspaces, and let $\hat{\mathcal{X}}$ be the hyperbolic space obtained from \mathcal{X} by coning off each $P \in \mathcal{P}$. Then we can endow \mathcal{X} with a hierarchical space structure as follows:

- the index-set \mathfrak{S} consists of \mathcal{P} together with an additional index S ;
- for all $P, Q \in \mathcal{P}$, we have $P \triangleleft Q$, while $P \not\sqsubseteq S$ for all $P \in \mathcal{P}$ (the orthogonality relation is empty and there is no other nesting);
- for each $P \in \mathcal{P}$, we let $\mathcal{C}P = P$;
- we declare $\mathcal{C}S = \hat{\mathcal{X}}$;
- the projection $\pi_S : \mathcal{X} \rightarrow \hat{\mathcal{X}}$ is the inclusion;
- for each $P \in \mathcal{P}$, let $\pi_P : \mathcal{X} \rightarrow P$ be the closest-point projection onto P (which is surjective);
- for each $P \in \mathcal{P}$, let ρ_S^P be the cone-point in $\hat{\mathcal{X}}$ associated to P ;
- for each $P \in \mathcal{P}$, let $\rho_P^S : \hat{\mathcal{X}} \rightarrow P$ be defined by $\rho_P^S(x) = \pi_P(x)$ for $x \in \mathcal{X}$, while $\rho_P^S(x) = \pi_P(Q)$ whenever x lies in the cone on $Q \in \mathcal{P}$.
- for distinct $P, Q \in \mathcal{P}$, let $\rho_Q^P = \pi_Q(P)$ (which is uniformly bounded since \mathcal{X} is hyperbolic relative to \mathcal{P}).

The above definition yields:

Theorem 9.3. *Let the geodesic metric space \mathcal{X} be hyperbolic relative to the collection \mathcal{P} of subspaces. Then, with \mathfrak{S} as above, we have that $(\mathcal{X}, \mathfrak{S})$ is a hierarchical space, and is moreover relatively hierarchically hyperbolic.*

Proof. By definition, for each $U \in \mathfrak{S}$, we have that either $U = S$ and $\mathcal{C}S = \hat{\mathcal{X}}$ is hyperbolic, or U is \sqsubseteq -minimal. The rest of the conditions of Definition 1.1 are verified as in the proof of Theorem 9.1. □

10. HIERARCHICAL HYPERBOLICITY OF 3-MANIFOLD GROUPS

In this section we show that fundamental groups of most 3-manifolds admit hierarchical hyperbolic structures. More precisely, we prove:

Theorem 10.1 (3-manifolds are hierarchically hyperbolic). *Let M be a closed 3-manifold. Then $\pi_1(M)$ is a hierarchically hyperbolic space if and only if M does not have a Sol or Nil component in its prime decomposition.*

Proof. It is well known that for a closed irreducible 3-manifold N the Dehn function of $\pi_1(N)$ is linear if N is hyperbolic, cubic if N is Nil, exponential if N is Sol, and quadratic in all other cases. Hence by Corollary 7.5, if $\pi_1(M)$ is a hierarchically hyperbolic space, then M does not contain Nil or Sol manifolds in its prime decomposition. It remains to prove the converse.

Since the fundamental group of any reducible 3-manifold is the free product of irreducible ones, the reducible case immediately follows from the irreducible case by Theorem 9.1.

When M is geometric and not Nil or Sol, then $\pi_1(M)$ is quasi-isometric to one of the following:

- \mathbb{R}^3 is hierarchically hyperbolic by Proposition 8.27;

- $\mathbb{H}^3, \mathbb{S}^3, \mathbb{S}^2 \times \mathbb{R}$ are (hierarchically) hyperbolic;
- $\mathbb{H}^2 \times \mathbb{R}$ and $PSL_2(\mathbb{R})$: the first is hierarchically hyperbolic by Proposition 8.27, whence the second is also since it is quasi-isometric to the first by [Rie01].

We may now assume M is not geometric. Our main step is to show that any irreducible non-geometric graph manifold group is a hierarchically hyperbolic space.

Let M be an irreducible non-geometric graph manifold. By [KL98, Theorem 2.3], by replacing M by a manifold whose fundamental group is quasi-isometric to that of M , we may assume that our manifold is a *flip graph manifold*, i.e., each Seifert fibered space component is a trivial circle bundles over a surfaces of genus at least 2 and each pair of adjacent Seifert fibered spaces are glued by flipping the base and fiber directions.

Let X be the universal cover of M . The decomposition of M into geometric components induces a decomposition of X into subspaces $\{S_v\}$, one for each vertex v of the Bass-Serre tree T of M . Each such subspace S_v is bi-Lipschitz homeomorphic to the product of a copy R_v of the real line with the universal cover Σ_v of a hyperbolic surface with totally geodesic boundary, and there are maps $\phi_v: S_v \rightarrow \Sigma_v$ and $\psi_v: S_v \rightarrow R_v$. Notice that Σ_v is hyperbolic, and in particular hierarchically hyperbolic. However, for later purposes, we endow Σ_v with the hierarchically hyperbolic structure originating from the fact that Σ_v is hyperbolic relative to its boundary components, see Theorem 9.1.

By Proposition 8.27 each S_v is a hierarchically hyperbolic space and thus we have a tree of hierarchically hyperbolic spaces. Each edge space is a product $\partial_0 \Sigma_v \times R_v$, where $\partial_0 \Sigma_v$ is a particular boundary component of Σ_v determined by the adjacent vertex. Further, since the graph manifold is flip, we also have that for each vertices v, w of the tree, the edge-hieromorphism between S_v and S_w sends $\partial_0 \Sigma_v$ to R_w and R_v to $\partial_0 \Sigma_w$.

We now verify the hypotheses of Theorem 8.6. The first hypothesis is that there exists k so that each edge-hieromorphism is k -hierarchically quasiconvex. This is easily seen since the edge-hieromorphisms have the simple form described above. The second hypothesis of Theorem 8.6, fullness of edge-hieromorphisms, also follows immediately from the explicit description of the edges here and the simple hierarchically hyperbolic structure of the edge spaces.

The third hypothesis of Theorem 8.6 requires that the tree has bounded supports. We can assume that the product regions S_v are maximal in the sense that each edge-hieromorphism sends the fiber direction R_v to $\partial_0 \Sigma_w$ in each adjacent S_w . It follows that the support of each \sim -class (in the language of Theorem 8.6) consists of at most 2 vertices. The last hypothesis of Theorem 8.6 is about non-orthogonality of maximal elements and again follows directly from the explicit hierarchically hyperbolic structure. Moreover, the part of the hypothesis about edge-spaces coinciding coarsely with standard product regions in vertex spaces follows from the explicit hierarchically hyperbolic structure.

All the hypotheses of Theorem 8.6 are satisfied, so $\pi_1 M$ (with any word metric) is a hierarchically hyperbolic space for all irreducible non-geometric graph manifolds M .

The general case that the fundamental group of any non-geometric 3-manifold is an hierarchically hyperbolic space now follows immediately by Theorem 9.1, since any 3-manifold group is hyperbolic relative to its maximal graph manifold subgroups. \square

Remark 10.2 ((Non)existence of HHG structures for 3-manifold groups). The proof of Theorem 10.1 shows that for many 3-manifolds M , the group $\pi_1 M$ is not merely a hierarchically hyperbolic space (when endowed with the word metric arising from a finite generating set), but is actually a hierarchically hyperbolic group. Specifically, if M is virtually compact special, then $\pi_1 M$ acts freely and cocompactly on a CAT(0) cube complex \mathcal{X} that is the universal cover of a compact special cube complex. Hence \mathcal{X} contains a $\pi_1 M$ -invariant factor system (see [BHS17a, Section 8]) consisting of a $\pi_1 M$ -finite set of convex subcomplexes.

This yields a hierarchically hyperbolic structure $(\mathcal{X}, \mathfrak{S})$ where $\pi_1 M \leq \text{Aut}(\mathfrak{S})$ acts cofinitely on \mathfrak{S} and geometrically on \mathcal{X} , i.e., $\pi_1 M$ is a hierarchically hyperbolic group.

The situation is quite different when $\pi_1 M$ is not virtually **compact** special. Indeed, when M is a nonpositively-curved graph manifold, $\pi_1 M$ virtually acts freely, but not necessarily cocompactly, on a CAT(0) cube complex \mathcal{X} , and the quotient is virtually special; this is a result of Liu [Liu13] which was also shown to hold in the case where M has nonempty boundary by Przytycki and Wise [PW14]. Moreover, $\pi_1 M$ acts with finitely many orbits of hyperplanes. Hence the $\pi_1 M$ -invariant factor system on \mathcal{X} from [BHS17a] yields a $\pi_1 M$ -equivariant HHS structure $(\mathcal{X}, \mathfrak{S})$ with \mathfrak{S} $\pi_1 M$ -finite. However, this yields an HHG structure on $\pi_1 M$ only if the action on \mathcal{X} is cocompact. In [HP15], the second author and Przytycki showed that $\pi_1 M$ virtually acts freely and cocompactly on a CAT(0) cube complex, with special quotient, only in the very particular situation where M is *chargeless*. This essentially asks that the construction of the hierarchically hyperbolic structure on \widetilde{M} from the proof of Theorem 10.1 can be done $\pi_1 M$ -equivariantly. In general, this is impossible: recall that we passed from \widetilde{M} to the universal cover of a flip manifold using a (nonequivariant) quasi-isometry. Motivated by this observation and the fact that the range of possible HHS structures on the universal cover of a JSJ torus is very limited, we conjecture that $\pi_1 M$ is a hierarchically hyperbolic group if and only if $\pi_1 M$ acts freely and cocompactly on a CAT(0) cube complex.

11. A NEW PROOF OF THE DISTANCE FORMULA FOR MAPPING CLASS GROUPS

We now describe the hierarchically hyperbolic structure of mapping class groups. In [BHS17a] we gave a proof of this result using several of the main results of [Beh06, BKMM12, MM99, MM00]. Here we give an elementary proof which is independent of the Masur–Minsky “hierarchy machinery.” One consequence of this is a new and concise proof of the celebrated Masur–Minsky distance formula [MM00, Theorem 6.12], which we obtain by combining Theorem 11.1 and Theorem 4.5.

- (1) Let S be closed connected oriented surface of finite type and let $\mathcal{M}(S)$ be its marking complex.
- (2) Let \mathfrak{S} be the collection of isotopy classes of essential subsurfaces of S , and for each $U \in \mathfrak{S}$ let $\mathcal{C}U$ be its curve complex.
- (3) The relation \sqsubseteq is nesting, \perp is disjointness and \pitchfork is overlapping.
- (4) For each $U \in \mathfrak{S}$, let $\pi_U: \mathcal{M}(S) \rightarrow \mathcal{C}U$ be the (usual) subsurface projection. For $U, V \in \mathfrak{S}$ satisfying either $U \sqsubseteq V$ or $U \pitchfork V$, denote $\rho_V^U = \pi_V(\partial U) \in \mathcal{C}V$, while for $V \sqsubseteq U$ let $\rho_V^U: \mathcal{C}U \rightarrow 2^{\mathcal{C}V}$ be the subsurface projection.

Theorem 11.1. *Let S be closed connected oriented surface of finite type. Then, $(\mathcal{M}(S), \mathfrak{S})$ is a hierarchically hyperbolic space, for \mathfrak{S} as above. In particular the mapping class group $\text{MCG}(S)$ is a hierarchically hyperbolic group.*

Proof. Hyperbolicity of curve graphs is the main result of [MM99]; more recent proofs of this were found in [Aou13, Bow14b, CRS15, HPW15, PS17], some of which are elementary.

Axioms 1, 2, 3 and 5 are clear (an elementary exposition of the Lipschitz condition for subsurface projections is provided in [MM00, Lemma 2.5], and the projections have quasi-convex image because they are coarsely surjective). Both parts of axiom 4 can be found in [Beh06]. The nesting part is elementary, and a short elementary proof in the overlapping case was obtained by Leininger and can be found in [Man10].

Axiom 7 was proven in [MM00], and an elementary proof is available in [Web15]. In fact, in the aforementioned papers it is proven that there exists a constant C so that for any subsurface W , markings x, y and geodesic from $\pi_W(x)$ to $\pi_W(y)$ the following holds. If

$V \sqsubseteq W$ and $V \neq W$ satisfies $d_V(x, y) \geq C$ then some curve along the given geodesic does not intersect ∂V . This implies Axiom 6, since we can take the T_i to be the complements of curves appearing along the aforementioned geodesic.

Axiom 8 follows easily from the following statement, which clearly holds: For any given collection of disjoint subsurfaces and curves on the given subsurfaces, there exists a marking on S that contains the given curves as base curves (or, up to bounded error, transversals in the case that the corresponding subsurface is an annulus).

Axiom 9 is hence the only delicate one. We are finished modulo this last axiom which we verify below in Proposition 11.2 (see also [BBF15, Proposition 5.11]). \square

Proposition 11.2. *($\mathcal{M}(S), \mathfrak{S}$) satisfies the uniqueness axiom, i.e., for each $\kappa \geq 0$, there exists $\theta_u = \theta_u(\kappa)$ such that if $x, y \in \mathcal{M}(S)$ satisfy $d_U(x, y) \leq \kappa$ for each $U \in \mathfrak{S}$ then $d_{\mathcal{M}(S)}(x, y) \leq \theta_u$.*

Proof. Note that when the complexity (as measured by the quantity $3g + p - 3$ where g is the genus and p the number of punctures) is less than 2 then $\mathcal{M}(S)$ is hyperbolic and thus the axiom holds. We will proceed by inducting on complexity: thus we will fix S to have complexity at least 2 and assume that all the axioms for a hierarchically hyperbolic space, including the uniqueness axiom, hold for each proper subsurface of S .

Now, having fixed our surface S , the proof is by induction on $d_{\mathcal{C}S}(base(x), base(y))$.

If $d_{\mathcal{C}S}(base(x), base(y)) = 0$, then x and y share some non-empty multicurve $\sigma = c_1 \cup \dots \cup c_k$. For x', y' the restrictions of x, y to $S - \sigma$ we have that, by induction, $d_{\mathcal{M}(S - \sigma)}(x', y')$ is bounded in terms of κ . We then take the markings in a geodesic in $\mathcal{M}(S - \sigma)$ from x' to y' and extend these all in the same way to obtain markings in $\mathcal{M}(S)$ which yield a path in $\mathcal{M}(S)$ from x to \hat{y} whose length is bounded in terms of κ , where \hat{y} is the marking for which:

- \hat{y} has the same base curves as y ,
- the transversal for each c_i is the same as the corresponding transversal for x , and
- the transversal for each curve in $base(y) - \{c_i\}$ is the same as the corresponding transversal for y .

Finally, it is readily seen that $d_{\mathcal{M}(S)}(\hat{y}, y)$ is bounded in terms of κ because the transversals of each c_i in the markings x and y are within distance κ of each other. This completes the proof of the base case of the Proposition.

Suppose now that that the statement holds whenever $d_{\mathcal{C}S}(base(x), base(y)) \leq n$, and let us prove it in the case $d_{\mathcal{C}S}(base(x), base(y)) = n + 1$. Let $c_x \in base(x)$ and $c_y \in base(y)$ satisfy $d_{\mathcal{C}S}(c_x, c_y) = n + 1$. Let $c_x = \sigma_0, \dots, \sigma_{n+1} = c_y$ be a tight geodesic (hence, each σ_i is a multicurve). Let σ be the union of σ_0 and σ_1 . Using the realization theorem in the subsurface $S - \sigma$ we can find a marking x' in $S - \sigma$ whose projections onto each $\mathcal{C}U$ for $U \subseteq S - \sigma$ coarsely coincide with $\pi_U(y)$. Let \hat{x} be the marking for which:

- $base(\hat{x})$ is the union of $base(x')$ and σ ,
- the transversal in \hat{x} of curves in $base(\hat{x}) \cap base(x')$ are the same as those in x' ,
- the transversal of c_x in \hat{x} is the same as the one in x ,
- the transversal in \hat{x} of a curve c in σ_1 is $\pi_{A_c}(y)$, where A_c is an annulus around c .

Note that $d_{\mathcal{C}S}(base(\hat{x}), base(y)) = n$. Hence, the following two claims conclude the proof.

Claim 1. $d_{\mathcal{M}(S)}(x, \hat{x})$ is bounded in terms of κ .

Proof. It suffices to bound $d_{\mathcal{C}U}(x, \hat{x})$ in terms of κ for each $U \subseteq S - c_x$. In fact, once we do that, by induction on complexity we know that we can bound $d_{\mathcal{M}(S - c_x)}(z, \hat{z})$, where z, \hat{z} are the restrictions of x, \hat{x} to $S - c_x$, whence the conclusion easily follows.

If U is contained in $S - \sigma$, then the required bound follows since $\pi_U(\hat{x})$ coarsely coincides with $\pi_U(x')$ in this case.

If instead ∂U intersects σ_1 , then $\pi_U(\hat{x})$ coarsely coincides with $\pi_U(\sigma_1)$.

At this point, we only have to show that $\pi_U(\sigma_1)$ coarsely coincides with $\pi_U(y)$, and in order to do so we observe that we can apply the bounded geodesic image theorem to the geodesic $\sigma_1, \dots, \sigma_{n+1}$. In fact, σ_1 intersects ∂U by hypothesis and σ_i intersects ∂U for $i \geq 3$ because of the following estimate that holds for any given boundary component c of ∂U

$$d_{C(S)}(\sigma_i, c) \geq d_{C(S)}(\sigma_i, \sigma_0) - d_{C(S)}(\sigma_0, c) \geq i - 1 > 1.$$

Finally, σ_2 intersects ∂U because of the definition of tightness: ∂U intersects σ_1 , hence it must intersect $\sigma_0 \cup \sigma_2$. However, it does not intersect σ_0 , whence it intersects σ_2 . ■

Claim 2. There exists κ' , depending on κ , so that for each subsurface U of S we have $d_{CU}(\hat{x}, y) \leq \kappa'$.

Proof. If σ_0 intersects ∂U , then $\pi_U(\hat{x})$ coarsely coincides with $\pi_U(\sigma_0)$. In turn, $\pi_U(\sigma_0)$ coarsely coincides with $\pi_U(x)$, which is κ -close to $\pi_U(y)$.

On the other hand, if U does not intersect σ , then we are done by the definition of x' .

Hence, we can assume that U is contained in $S - \sigma_0$ and that σ_1 intersects ∂U . In particular, $\pi_U(\hat{x})$ coarsely coincides with $\pi_U(\sigma_1)$. But we showed in the last paragraph of the proof of Claim 1 that $\pi_U(\sigma_1)$ coarsely coincides with $\pi_U(y)$, so we are done. ■

As explained above, the proofs of the above two claims complete the proof. □

REFERENCES

- [AB] Carolyn Abbott and Jason Behrstock. Conjugator lengths in hierarchically hyperbolic groups. *preprint arXiv:1808.09604*.
- [ABD] Carolyn Abbott, Jason Behrstock, and Matthew G Durham. Largest acylindrical actions and stability in hierarchically hyperbolic groups. *preprint arXiv:1705.06219*.
- [Ago13] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.
- [Aou13] Tarik Aougab. Uniform hyperbolicity of the graphs of curves. *Geometry & Topology*, 17(5):2855–2875, 2013.
- [BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publications mathématiques de l’IHÉS*, 122(1):1–64, 2015.
- [BDS11] J. Behrstock, C. Druţu, and M. Sapir. Median structures on asymptotic cones and homomorphisms into mapping class groups. *Proc. London Math. Soc.*, 102(3):503–554, 2011.
- [Beh06] J. Behrstock. Asymptotic geometry of the mapping class group and Teichmüller space. *Geometry & Topology*, 10:2001–2056, 2006.
- [BF92] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geometry*, 35:85–101, 1992.
- [BF02] M. Bestvina and K. Fujiwara. Bounded cohomology of subgroups of mapping class groups. *Geometry & Topology*, 6:69–89 (electronic), 2002.
- [BF14a] Mladen Bestvina and Mark Feighn. Hyperbolicity of the complex of free factors. *Adv. Math.*, 256:104–155, 2014.
- [BF14b] Mladen Bestvina and Mark Feighn. Subfactor projections. *J. Topol.*, 7(3):771–804, 2014.
- [BH16] J. Behrstock and M.F. Hagen. Cubulated groups: thickness, relative hyperbolicity, and simplicial boundaries. *Geometry, Groups, and Dynamics*, 10(2):649–707, 2016.
- [BHS17a] J. Behrstock, M.F. Hagen, and A. Sisto. Hierarchically hyperbolic spaces I: curve complexes for cubical groups. *Geom. Topol.*, 21:1731–1804, 2017.
- [BHS17b] Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proceedings of the London Mathematical Society*, 114(5):890–926, 2017.
- [BHS17c] Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *preprint arXiv:1704.04271*, 2017.

- [BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. *Geom. Topol.*, 16(2):781–888, 2012.
- [BM08] Jason A Behrstock and Yair N Minsky. Dimension and rank for mapping class groups. *Annals of Mathematics*, 167:1055–1077, 2008.
- [BM11] Jason A. Behrstock and Yair N. Minsky. Centroids and the rapid decay property in mapping class groups. *J. Lond. Math. Soc. (2)*, 84(3):765–784, 2011.
- [Bow13] Brian H. Bowditch. Coarse median spaces and groups. *Pacific J. Math*, 261(1):53–93, 2013.
- [Bow14a] Brian H Bowditch. Embedding median algebras in products of trees. *Geometriae Dedicata*, 170(1):157–176, 2014.
- [Bow14b] Brian H. Bowditch. Uniform hyperbolicity of the curve graphs. *Pacific J. Math*, 269(2):269–280, 2014.
- [Bow18] Brian H. Bowditch. Large-scale rigidity properties of the mapping class groups. *Pacific J. Math.*, 293(1):1–73, 2018.
- [BR18] Federico Berlai and Bruno Robbio. A refined combination theorem for hierarchically hyperbolic groups. *In preparation*, 2018.
- [Bri00] Peter Brinkmann. Hyperbolic automorphisms of free groups. *Geometric & Functional Analysis*, 10(5):1071–1089, 2000.
- [Bri10] Martin R Bridson. Semisimple actions of mapping class groups on CAT(0) spaces. *Geometry of Riemann surfaces*, 368:1–14, 2010.
- [BS05] S Buyalo and P Svetlov. Topological and geometric properties of graph-manifolds. *St. Petersburg Mathematical Journal*, 16(2):297–340, 2005.
- [BV] M R Bridson and K Vogtmann. The Dehn functions of $Out(F_n)$ and $Aut(F_n)$. *arXiv:1011.1506*.
- [BV95] M R Bridson and K Vogtmann. On the geometry of the automorphism group of a free group. *Bulletin of the London Mathematical Society*, 27(6):544–552, 1995.
- [CDH10] Indira Chatterji, Cornelia Druțu, and Frédéric Haglund. Kazhdan and Haagerup properties from the median viewpoint. *Advances in Mathematics*, 225(2):882–921, 2010.
- [Che00] Victor Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.
- [CLM12] Matt T. Clay, Christopher J. Leininger, and Johanna Mangahas. The geometry of right-angled Artin subgroups of mapping class groups. *Groups Geom. Dyn.*, 6(2):249–278, 2012.
- [CRS15] Matt Clay, Kasra Rafi, and Saul Schleimer. Uniform hyperbolicity of the curve graph via surgery sequences. *Algebraic and Geometric Topology*, 14(6):3325–3344, 2015.
- [DHS17] Matthew Durham, Mark Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.
- [Dil50] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Ann. of Math*, 51:161–166, 1950.
- [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.
- [DT15] Matthew Gentry Durham and Samuel J. Taylor. Convex cocompactness and stability in mapping class groups. *Algebr. Geom. Topol.*, 15(5):2839–2859, 2015.
- [FLS15] Roberto Frigerio, Jean-François Lafont, and Alessandro Sisto. Rigidity of high dimensional graph manifolds. *Astérisque*, (372):xxi+177, 2015.
- [Ger92] Stephen M Gersten. *Dehn functions and ℓ_1 -norms of finite presentations*. Springer, 1992.
- [GHR03] Steve M Gersten, Derek F Holt, and Tim R Riley. Isoperimetric inequalities for nilpotent groups. *Geometric & Functional Analysis*, 13(4):795–814, 2003.
- [Hae16] Thomas Haettel. Higher rank lattices are not coarse median. *Algebr. Geom. Topol.*, 16(5):2895–2910, 2016.
- [Hag07] Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. *arXiv:0705.3386*, 2007.
- [Hag14] Mark F. Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.*, 7(2):385–418, 2014.
- [HM13a] Michael Handel and Lee Mosher. The free splitting complex of a free group, I: hyperbolicity. *Geom. Topol.*, 17(3):1581–1672, 2013.
- [HM13b] Michael Handel and Lee Mosher. Lipschitz retraction and distortion for subgroups of $Out(F_n)$. *Geometry & Topology*, 17(3):1535–1579, 2013.
- [HP15] Mark F. Hagen and Piotr Przytycki. Cocompactly cubulated graph manifolds. *Israel J. Math.*, 207(1):377–394, 2015.
- [HPW15] Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb. 1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs. *J. Eur. Math. Soc. (JEMS)*, 17(4):755–762, 2015.
- [HS16] Mark F Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. *preprint arXiv:1609.01313v2*, 2016.

- [HW15] Mark F. Hagen and Daniel T. Wise. Cubulating hyperbolic free-by-cyclic groups: the general case. *Geometric and Functional Analysis*, 25(1):134–179, 2015.
- [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.
- [KL96] Michael Kapovich and Bernhard Leeb. Actions of discrete groups on nonpositively curved spaces. *Math. Ann.*, 306(2):341–352, 1996.
- [KL97] M. Kapovich and B. Leeb. Quasi-isometries preserve the geometric decomposition of Haken manifolds. *Invent. Math.*, 128(2):393–416, 1997.
- [KL98] M. Kapovich and B. Leeb. 3-manifold groups and nonpositive curvature. *Geom. Funct. Anal.*, 8(5):841–852, 1998.
- [Liu13] Yi Liu. Virtual cubulation of nonpositively curved graph manifolds. *J. Topol.*, 6(4):793–822, 2013.
- [Man10] Johanna Mangahas. Uniform exponential growth of subgroups of the mapping class group. *Geometric and Functional Analysis*, 19(5):1468–1480, 2010.
- [MM99] Howard A Masur and Yair N Minsky. Geometry of the complex of curves I: Hyperbolicity. *Inventiones mathematicae*, 138(1):103–149, 1999.
- [MM00] Howard A Masur and Yair N Minsky. Geometry of the complex of curves II: Hierarchical structure. *Geometric and Functional Analysis*, 10(4):902–974, 2000.
- [MO14] Ashot Minasyan and Denis Osin. Acylindrical hyperbolicity of groups acting on trees. *Mathematische Annalen*, pages 1–51, 2014.
- [Mos95] L. Mosher. Mapping class groups are automatic. *Ann. of Math.*, 142:303–384, 1995.
- [Mou17] Sarah C Mousley. Exotic limit sets of Teichmüller geodesics in the hhs boundary. *preprint arXiv:1704.08645*, 2017.
- [Mou18] Sarah C. Mousley. Nonexistence of boundary maps for some hierarchically hyperbolic spaces. *Algebr. Geom. Topol.*, 18(1):409–439, 2018.
- [MR18] Sarah C Mousley and Jacob Russell. Hierarchically hyperbolic groups are determined by their morse boundaries. *preprint arXiv:1801.04867*, 2018.
- [PS17] Piotr Przytycki and Alessandro Sisto. A note on acylindrical hyperbolicity of mapping class groups. In *Hyperbolic geometry and geometric group theory*, volume 73 of *Adv. Stud. Pure Math.*, pages 255–264. Math. Soc. Japan, Tokyo, 2017.
- [PW14] Piotr Przytycki and Daniel T. Wise. Graph manifolds with boundary are virtually special. *J. Topol.*, 7(2):419–435, 2014.
- [Rie01] E. Rieffel. Groups quasi-isometric to $\mathbf{H}^2 \times \mathbf{R}$. *J. London Math. Soc. (2)*, 64(1):44–60, 2001.
- [Sis11] Alessandro Sisto. 3-manifold groups have unique asymptotic cones. *arXiv:1109.4674*, 2011.
- [Sis12] Alessandro Sisto. On metric relative hyperbolicity. *arXiv:1210.8081*, 2012.
- [Sis13] Alessandro Sisto. Projections and relative hyperbolicity. *Enseign. Math. (2)*, 59(1-2):165–181, 2013.
- [Spr17] Davide Spriano. Hyperbolic HHS I: Factor systems and quasi-convex subgroups. *preprint arXiv:1711.10931*, 2017.
- [Spr18] Davide Spriano. Hyperbolic HHS II: Graphs of hierarchically hyperbolic groups. *preprint arXiv:1801.01850*, 2018.
- [SS] L. Sabalka and D. Savchuk. Submanifold projection. *ARXIV:1211.3111*.
- [Vok17] Kate M Vokes. Hierarchical hyperbolicity of the separating curve graph. *arXiv preprint arXiv:1711.03080*, 2017.
- [Web15] Richard C. H. Webb. Uniform bounds for bounded geodesic image theorems. *J. Reine Angew. Math.*, 709:219–228, 2015.
- [Wis14] Daniel Wise. Cubular tubular groups. *Transactions of the American Mathematical Society*, 366(10):5503–5521, 2014.
- [Woo16] Daniel J. Woodhouse. Classifying finite dimensional cubulations of tubular groups. *Michigan Math. J.*, 65(3):511–532, 2016.

BARNARD COLLEGE, COLUMBIA UNIVERSITY, NEW YORK, NEW YORK, USA

Current address: Lehman College and The Graduate Center, CUNY, New York, New York, USA

E-mail address: `jason@math.columbia.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`