

HIERARCHICAL HYPERBOLICITY, BALL QUOTIENTS, AND THE MODULI SPACE OF SMOOTH CUBIC SURFACES

DANIEL ALLCOCK AND ALEX WRIGHT

ABSTRACT. We prove that the orbifold fundamental group of the moduli space of smooth complex cubic surfaces is a hierarchically hyperbolic group. This places it within a powerful geometric framework whose prototypical examples are mapping class groups. The proof uses a known incomplete complex hyperbolic metric on the moduli space, and applies in a broader setting which includes a number of other moduli spaces.

CONTENTS

1. Introduction	2
2. Background on geometry	8
3. Criteria for hierarchical hyperbolicity	12
4. The model space and statement of the goal	14
5. Stratum domains and the branched cover	16
6. Swaddling and curvature	20
7. Cusp domains	26
8. Angle domains	35
9. Nesting and orthogonality	43
10. Active regions	48
11. Proof of hierarchical hyperbolicity	53
Appendix A. Proof of criteria for hierarchical hyperbolicity	60
Appendix B. The combinatorics of the index set (not required)	64
Appendix C. Relative hierarchical hyperbolicity (not required)	69
Appendix D. Other examples (not required)	72
References	77

1. INTRODUCTION

1.1. Main results. The moduli space $\mathcal{M}_{g,k}$ of Riemann surfaces with genus g and k punctures dates from the 19th century. The corresponding mapping class group is its orbifold fundamental group, and has been a focus of geometric group theory for decades. Previously most of our understanding of these groups relied on constructions specific to them, but starting about ten years ago they and several other major classes of groups have been fruitfully studied from the single perspective of hierarchically hyperbolic groups. In this paper we show that this theory also applies to many groups from algebraic geometry, including the fundamental group of another 19th century moduli space:

Theorem 1.1. *The orbifold fundamental group $\pi_1(\mathcal{M}_{cs})$, of the moduli space $\mathcal{M}_{cs}$ of smooth complex cubic surfaces, is a hierarchically hyperbolic group (HHG).*

Briefly, cubic surfaces are the subvarieties of $\mathbb{CP}^3$ defined by homogeneous cubic polynomials, and $\mathcal{M}_{cs}$ is the space (orbifold) that parameterizes the smooth ones up to isomorphism (or equivalently projective equivalence). Our hierarchical hyperbolic structure was inspired by a structural similarity between $\mathcal{M}_{cs}$ and $\mathcal{M}_{g,k}$.

Our proof does not use any properties of cubic surfaces. It relies on the fact that $\mathcal{M}_{cs}$ can be described as the complex 4-ball (also called complex hyperbolic 4-space), minus an arrangement of hyperplanes with certain properties, modulo a lattice in $PU(4,1)$ [ACT02a]. Specifically, Theorem 1.1 follows from our next result. See Section 11.4 for this implication and Appendix D for many more examples.

Theorem 1.2 (Main theorem). *Suppose $\mathcal{H}$ is a union of hyperplanes in the complex n -ball $B := B^n$ satisfying the following.*

- (A) *Each point of B has a neighborhood meeting only finitely many hyperplanes.*
- (B) *Any two hyperplanes that meet do so orthogonally.*
- (C) *$\mathcal{H}$ is invariant under a lattice $\Gamma \subseteq PU(n,1)$.*
- (D) *Each cusp of Γ is a limit point of a 1-dimensional stratum.*

Then the orbifold fundamental group $\hat{\Gamma}$ of the complex analytic orbifold $\Gamma \backslash (B - \mathcal{H})$ is hierarchically hyperbolic.

We restrict the term “hyperplane in B ” to the hyperplanes comprising $\mathcal{H}$, because we will refer to them often and never to other hyperplanes in B . The stratification in (D) is the one induced by any hyperplane arrangement satisfying (A)–(B): the codimension k strata are the connected components of the set of points contained in exactly k hyperplanes. Whenever we use fundamental groups and universal covers it will be in the orbifold sense. For example, $B - \mathcal{H}$ is a covering space of $\Gamma \backslash (B - \mathcal{H})$, so $\hat{\Gamma}$ contains $\pi_1(B - \mathcal{H})$ as a normal subgroup, with quotient Γ .

We will expand on what an HHG is, but first mention some consequences: solvability of the word and conjugacy problems ([BHS19, Corollary 7.5] and [HHP23, Corollary H]), analogues of the Nielsen-Thurston classification and the Tits alternative ([DHS17, Proposition 6.18, Theorem 7.1] and [DHS20, Theorem 4.1]), and more (see below). All of these are new for $\pi_1(\mathcal{M}_{cs})$.

The lattice Γ is allowed to be non-uniform, in which case the theorem stands in sharp contrast with the fact that non-uniform lattices Γ in $PU(n > 1, 1)$ are never HHGs. This is the (known) $\mathcal{H} = \emptyset$ case of Remark C.4: if (A), (B) and (C) hold

but (D) does not, then $\hat{\Gamma}$ cannot be an HHG. But it is a relatively hierarchically hyperbolic group (RHHG) by [Theorem C.1](#).

Other examples. A remarkable number of spaces have the form $\Gamma \backslash (B - \mathcal{H})$; see for example [\[Loo16\]](#) and [\[DK07\]](#). [Appendix D](#) lists other examples to which [Theorem 1.2](#) or [Theorem C.1](#) applies, including:

- the moduli space of stable cubic threefolds;
- certain moduli spaces of K3 surfaces with non-symplectic automorphisms;
- certain moduli spaces of abelian varieties with extra structure, which are open subsets of certain unitary Shimura varieties of PEL type; and
- the moduli spaces $\mathcal{M}_{0,5}/S_5$ and $\mathcal{M}_{0,6}/S_6$ of 5 resp. 6 points in $\mathbb{CP}^1$.

1.2. Hierarchical hyperbolicity. We pause now to introduce hierarchical hyperbolicity for those not familiar with it.

What is a hierarchically hyperbolic group? A good part of the study of mapping class groups has been built on deep breakthrough results of Masur-Minsky [\[MM99, MM00\]](#) relating to the geometry of curve complexes. These complexes encode how the strata of the Deligne-Mumford compactification $\overline{\mathcal{M}}_{g,k}$ fit together, after appropriately lifting to the universal cover of $\mathcal{M}_{g,k}$.

After much work by many people developing and applying the resulting theory, Behrstock-Hagen-Sisto found strong parallels elsewhere in geometric group theory, and introduced the definition of a hierarchically hyperbolic group [\[BHS17b, BHS19\]](#). By design, the restrictive axioms of hierarchical hyperbolicity guarantee that many results on mapping class groups apply to all hierarchically hyperbolic groups. Fresh excitement and having the right category to work in have facilitated many further results, and many new examples have been found within geometric group theory and low dimensional topology.

This historical discussion gives us our first answer to the question:

- (1) A hierarchically hyperbolic group is one that is geometrically similar enough to a mapping class group that many results about mapping class groups carry over.

There are other equally high-level answers:

- (2) A hierarchically hyperbolic group is one that admits a system of coordinates with highly structured redundancy, such that the coordinates are valued in (Gromov) hyperbolic spaces, and the different coordinates give information on the geometry of different parts of the group and how they fit together.
- (3) A hierarchically hyperbolic group is one that looks hyperbolic, except for product regions whose structure and overlaps are similarly constrained.

The precise definition of hierarchical hyperbolicity is intricate, but it is not truly required for our work. We state a sufficient criterion for hierarchical hyperbolicity in [Section 3](#), deferring the definition of hierarchical hyperbolicity and the proof of the criterion to [Appendix A](#).

See [\[Wri26\]](#) for a brief introduction to hierarchical hyperbolicity, [\[Sis19\]](#) for a deeper introduction, and [\[Sis25\]](#) for a more advanced and up to date survey.

Other consequences of hierarchical hyperbolicity. We think of [Theorem 1.2](#) not as supplying a few isolated consequences but rather a coherent geometric theory.

To illustrate the far reaching nature of this theory, we briefly discuss the Farrell-Jones Conjecture. This conjecture, originating in [FJ93], concerns the algebraic K-theory and L-theory of group rings and has strong implications for the Borel and Novikov Conjectures; see for example the ICM proceedings [Lüc10, Bar18]. The Farrell-Jones Conjecture was proven for $\pi_1(\mathcal{M}_{g,k})$ in [BB19], and for large classes of HHGs in [DMS25]. In Section 11.5 we use [DMS25, Theorem C] to conclude the following.

Theorem 1.3. *The Farrell-Jones Conjecture holds for all $\hat{\Gamma}$ as in Theorem 1.2.*

Other results on hierarchically hyperbolic groups, too numerous to list here, are proven for example in [BHS17b, BHS19, BHS17a, DHS17, ABD21, BHS21, PS23, RST23, HHP23, DMS23, Pet26, DMS25, AH26]. Some results on hierarchically hyperbolic groups require additional assumptions, and we show our $\hat{\Gamma}$ satisfy the most common such assumptions in Appendix B (note also Remarks 9.11 and 11.10).

1.3. The proof. We now wish to discuss the proof of Theorem 1.2 and the difficulties it must overcome.

The motivating example. But let us step back for a moment and first discuss the motivating example $\pi_1(\mathcal{M}_{cs})$. A naive comparison of $\pi_1(\mathcal{M}_{cs})$ to $\pi_1(\mathcal{M}_{g,k})$ gives rise to the hope that the topological monodromy

$$\pi_1(\mathcal{M}_{cs}) \rightarrow \pi_0(\text{Homeo}(\Sigma_{cs}))$$

is an isomorphism. Here Σ_{cs} is a cubic surface, and $\pi_0(\text{Homeo}(\Sigma_{cs}))$ is its topological mapping class group. However, $\pi_1(\mathcal{M}_{cs})$ in fact has *finite* image in the topological mapping class group $\pi_0(\text{Homeo}(\Sigma_{cs}))$; see for example [Shi26, Theorem 2.2] or [Har79, Section III.3] together with [Qui86, Theorem 1.1] and [GGH⁺26]. This means that any attempt to study $\pi_1(\mathcal{M}_{cs})$ via direct topological analogues of objects like curves on surfaces will find little traction.

The key tool. We will make an analogy between $\mathcal{M}_{g,k}$ and $\Gamma \backslash (B - \mathcal{H})$, by thinking of $\mathcal{M}_{g,k}$ as the quotient of its universal cover (Teichmüller space $\mathcal{T}_{g,k}$) by the mapping class group $\pi_1(\mathcal{M}_{g,k})$. The Weil-Petersson metric on $\mathcal{M}_{g,k}$ becomes analogous to the complex hyperbolic metric on $\Gamma \backslash (B - \mathcal{H})$.

The metric completion $\overline{\mathcal{T}}_{g,k}$ of $\mathcal{T}_{g,k}$ with respect to the lifted Weil-Petersson metric is known as augmented Teichmüller space or the Deligne-Mumford bordification. Its quotient by $\pi_1(\mathcal{M}_{g,k})$ is the Deligne-Mumford compactification $\overline{\mathcal{M}}_{g,k}$ [Mas76]. The boundary $\overline{\mathcal{T}}_{g,k} - \mathcal{T}_{g,k}$ is a union of (complex) codimension 1 components that are in bijection with (essential, non-peripheral) simple closed curves on a topological reference surface $\Sigma_{g,k}$. The curve complex of $\Sigma_{g,k}$ can be coarsely identified with a kind of quotient of $\overline{\mathcal{T}}_{g,k}$ called an electrification [MM99, Lemma 7.1]. Electrification collapses each of the boundary components to have finite diameter.

In the complex ball situation, our analogy suggests thinking of the universal cover of $\Gamma \backslash (B - \mathcal{H})$ as an analogue of $\mathcal{T}_{g,k}$. This is the same as the universal cover of $B - \mathcal{H}$, except with deck group $\hat{\Gamma}$ not just its subgroup $\pi_1(B - \mathcal{H})$. In the motivating example, this is the universal cover of the moduli space of smooth complex cubic surfaces. The metric completion $\hat{B}$ of the universal cover, with respect to the lifted complex hyperbolic metric, will be our key tool. It is a CAT(−1) space [All00a], obtained by adjoining to the universal cover an infinite family of (complex)

codimension 1 components which meet each other “the same way” as the boundary components of $\overline{\mathcal{T}}_{g,k}$. (This uses the orthogonality hypothesis of Theorem 1.2.)

Although there is no known way to use the Weil-Petersson metric to prove that $\pi_1(\mathcal{M}_{g,n})$ is an HHG, the analogy between $\hat{B}$ and $\overline{\mathcal{T}}_{g,k}$ provides motivation for our result that $\hat{\Gamma}$ is an HHG. This analogy is strongest when Γ is cocompact, but when Γ has cusps the geometry of $\hat{B}$ near its cusps (suitably defined) has new features.

The hyperbolic spaces. Since $\hat{B}$ is $\text{CAT}(-1)$ it is of course (Gromov) hyperbolic, so it may surprise non-experts that $\hat{B}$ cannot typically be used as one of the (Gromov) hyperbolic spaces required by the definition of hierarchical hyperbolicity. In fact we use four kinds of (Gromov) hyperbolic spaces:

- (1) Electrifications of stratum closures in $\hat{B}$, indexed by “stratum domains”. This includes an electrification of $\hat{B}$ itself, at the top of our hierarchy.
- (2) Spaces quasi-isometric to $\mathbb{R}$, which record a sort of angular coordinate analogous to a Fenchel-Nielsen twist parameter. We index these by “angle domains”, and there is one for each codimension 1 stratum of $\hat{B}$.
- (3) More spaces quasi-isometric to $\mathbb{R}$, which record the “vertical” direction as measured from each cusp, and are associated to the center of the Heisenberg group. We index these by “center domains”.
- (4) Spaces quasi-isometric to trees, which record the “horizontal” directions as measured from each cusp. We index these by “tree domains”, and in combination with center domains they witness how (D) transforms the cusp geometry of Γ from an obstruction to hierarchical hyperbolicity into an entirely different and more compatible geometry.

Swaddling. Suppose $\hat{S}$ is the closure of a (complex) codimension 1 boundary stratum in $\hat{B}$. One can make sense of its unit normal bundle, which is a principal $\mathbb{R}$ -bundle over a singular space. It has something like a connection, pulled back from the unit normal bundle of the hyperplane of B over which $\hat{S}$ lies. If this were flat, then it would be easy to define an $\mathbb{R}$ -valued function describing angular position around $\hat{S}$. But it is not. So we introduce a construction called “swaddling”, that starts with a principal $\mathbb{R}$ -bundle and “squeezes it horizontally”. This produces either a quasiline or a finite-diameter space, and it is here that the “angle function” around $\hat{S}$ takes values. One might view this as a geometric alternative to quasi-morphism based ideas introduced in the HHG context in [HRSS25] and used in [FFMS26, HMS24, DDLS24].

We find another use for swaddling when examining the cusps of $\hat{B}$. Consider a horosphere ∂T_c centered at a cusp of Γ ; it is the boundary of a horoball T_c . It is a copy of the Heisenberg group, and has the structure of a principal $\mathbb{R}$ -bundle over $\mathbb{C}^{n-1}$ with connection. A lift $\partial \hat{T}_c$ of ∂T_c to $\hat{B}$ is a singular space but still a principal $\mathbb{R}$ -bundle, which we can swaddle after pulling back the connection. Although swaddling ∂T_c yields a finite-diameter space, swaddling $\partial \hat{T}_c$ yields a quasiline and leads to our center domains. In fact, $\partial \hat{T}_c$ is very much like the product of one copy of $\mathbb{R}$ and $n - 1$ trees. Seeing the transition to this structure, starting from Heisenberg geometry, was a lovely moment.

To show that these swaddlings yield quasilines rather than finite-diameter spaces, we develop quantitative tools that allow us to show that a given set of swaddling data is “almost flat”. This is similar to checking that a 2-cocycle is bounded. In

one case we prove a local-to-global result that allows us to lift an easy boundedness result from B to a more subtle one on $\hat{B}$. This result should be applicable to many other swaddlings.

Architecture. We begin with background in [Section 2](#) and our (skippable until later) criterion for hierarchical hyperbolicity in [Section 3](#). [Section 4](#) through [Section 8](#) develop the geometry of the four kinds of domains, including swaddling and much more. We work extensively with geodesics in $\hat{B}$, because we lack tools like the curves on Riemann surfaces used for the mapping class groups. This work is complicated by the fact that $\hat{B}$ is not locally compact, and the fact that geodesics can leave the “thick part” $\hat{B}^\theta$, which is our geometric model for the group $\hat{\Gamma}$. We establish various properties of nearest-point projection maps, the hyperbolicity of the stratum domains, and the manner in which a horosphere in $\hat{B}$ is “like” a product of $\mathbb{R}$ and some trees. Unlike in some constructions of hierarchically hyperbolic groups, we cannot use tree or quasi-tree behavior in our analysis, since our stratum domains do not give quasi-trees. In [Section 9](#) through [Section 11](#) we must fit together the four very different types of domains to build the HHG structure on $\hat{\Gamma}$.

1.4. Previous work. Loosely speaking, $\Gamma \backslash B^\circ$ might be considered a version of $\Gamma \backslash B$ with $\Gamma \backslash \mathcal{H}$ drilled out, and thus our [Theorem 1.2](#) might be viewed as related to work like [\[GHM⁺24\]](#) on drilling hyperbolic groups and manifolds.

[\[Bel12a, Theorem 1.1\(ii\)\]](#) shows that if distinct hyperplanes of $\mathcal{H}$ are disjoint, and if $\Gamma \backslash \mathcal{H}$ is compact, then $\hat{\Gamma}$ is relatively hyperbolic. The proof uses a delicate warped product construction to build a complete metric on $\Gamma \backslash B^\circ$, and the $\hat{\Gamma}$ studied are of intrinsic interest in relation to rigidity phenomena and geometric finiteness. This line of research was continued in [\[BH13\]](#), and [\[BH13, Theorem 1.4\]](#) gives in particular that $\hat{\Gamma}$ is relatively hyperbolic under a “sparsity” condition. See also [\[Bel12b, Min18\]](#) for related results. We note in [Remark C.5](#) that very strong assumptions are required for one of our $\hat{\Gamma}$ to be relatively hyperbolic, and in particular $\pi_1(\mathcal{M}_{cs})$ is not relatively hyperbolic; so the sparsity condition of [\[BH13\]](#) cannot hold for our key example.

When $n > 1$ and $\mathcal{H} \neq \emptyset$, [\[BH13, Theorem 1.5\]](#) shows that our $\hat{\Gamma}$ cannot be isomorphic to a discrete isometry group of a Hadamard manifold of pinched negative curvature, and [\[ACT02b, Theorem 1.2\]](#) shows that it cannot be isomorphic to a lattice in a Lie group with finitely many connected components. [\[ACT02b, Theorem 1.1\]](#) shows that the infinitely generated kernel of $\hat{\Gamma} \rightarrow \Gamma$ has a presentation with a certain form, and [\[Loo08\]](#) and [\[ABL26\]](#) give simple finite presentations for $\pi_1(\mathcal{M}_{cs})$. A recent result shows that the divisor subgroup of $\pi_1(\mathcal{M}_{cs})$ is characteristic, and that $\mathcal{M}_{cs}$ has no non-trivial holomorphic automorphisms [\[BFJS26\]](#).

At a technical level our work can be compared to [\[HMS24\]](#), which proves that Artin groups of large and hyperbolic type are HHGs, like us starting with $\text{CAT}(-1)$ geometry. Unlike ours, their $\text{CAT}(-1)$ spaces are metrized 2-dimensional simplicial complexes.

1.5. Questions and conjectures. Our work might lead in many directions.

Beyond hierarchical hyperbolicity. For all that [Theorem 1.1](#) goes a long way to moving the state of knowledge on $\pi_1(\mathcal{M}_{cs})$ towards catching up with $\pi_1(\mathcal{M}_{g,k})$, this only applies for questions of a certain flavor. An important and natural question that our work says nothing about is the following.

Question 1.4. Is $\pi_1(\mathcal{M}_{cs})$ residually finite?

There are also questions for which [Theorem 1.1](#) may help a great deal, but which require additional work, such as automaticity and quasi-isometric rigidity.

Beyond orthogonal arrangements. [\[All13, Conjecture 1.4\]](#) and [\[PP16, Conjecture 1.1\]](#), if true, would allow much of our analysis to proceed for naturally occurring examples where our orthogonality assumption [\(B\)](#) does not hold. For example, the moduli space of smooth complex cubic threefolds has the form $\Gamma \backslash (B - \mathcal{H})$ with $\mathcal{H}$ non-orthogonal [\[All03, LS07\]](#).

Conjecture 1.5. The fundamental group of the moduli space of smooth complex cubic threefolds is an HHG.

We similarly conjecture that the $\hat{\Gamma}$ that appear in relation to Allcock’s monstrous proposal [\[All09\]](#) and in relation to moduli of rational elliptic surfaces [\[HL02\]](#) are (R)HHGs.

Beyond ball quotients. We next move on to the setting of type IV domains, where the complex ball is replaced by the $SO(2, n)$ hermitian symmetric space. Remarkably, some moduli spaces, such as $\mathcal{M}_{0,5}$ and $\mathcal{M}_{cs}$, fit in this setup as well as the ball quotient setup (see [Remark D.7](#) and [\[MSY92\]](#)), and such common examples are among the reasons to hope our techniques can be useful in the type IV setting.

Here it is likely that, although some examples are (R)HHGs, many others are not. As is the case for outer automorphism groups of free groups, it may be that even examples that are not hierarchically hyperbolic can be understood using ideas from hierarchical hyperbolicity.

Key examples of interest include the following; all are of the form $\Gamma \backslash (X - \mathcal{H})$ for X a type IV domain and $\mathcal{H}$ a (typically not orthogonal) hyperplane arrangement.

- (1) The moduli space of smooth complex Enriques surfaces, where $\mathcal{H}$ is orthogonal; see [\[Hor78a, Hor78b, Nam85, All00c, Dol16a, All00a\]](#).
- (2) The moduli space of cubic fourfolds; see [\[Voi86, Loo09, Laz10\]](#).
- (3) Moduli spaces of amply polarized K3 surfaces of degree $2d$, where we expect some of the behavior to depend on d ; see [\[Pvv71, Huy23, AE25, GHS13\]](#). And similarly in the more general setting of lattice-polarized K3 surfaces [\[Dol96\]](#).
- (4) Certain spaces that arise in singularity theory, especially related to unimodal singularities; see [\[All13, Section 8\]](#) and [\[Loo78, Loo84, Loo81, Laz09, Zha24\]](#).

A great deal of work has been devoted to compactifications of many of the spaces we have discussed, for example [\[Loo03a, Loo03b, HKT09, CMGHL24, Nar82, Loo16, AEH24, LO19, LO18, GKS21, FSW25, AF02, OSS16\]](#). Just as the Deligne-Mumford compactification of $\mathcal{M}_{g,k}$ is intimately related to the HHG structure on $\pi_1(\mathcal{M}_{g,k})$, some of these compactifications may be relevant to the geometric group theory of the relevant moduli spaces.

A symplectic mapping class group? Besides its interest in algebraic geometry, our motivating group $\pi_1(\mathcal{M}_{cs})$ arises at what might be a fertile threshold in the theory of symplectic mapping class groups. Using a Fubini-Study symplectic form ω_{FS} on a cubic surface Σ_{cs} , one gets a symplectic monodromy map

$$\pi_1(\mathcal{M}_{cs}) \rightarrow \pi_0(\mathrm{Symp}(\Sigma_{cs}, \omega_{FS})),$$

with the right side called the symplectic mapping class group (of cubic surfaces). When cubic surfaces are replaced by more-complicated varieties, one does not expect the corresponding maps to be isomorphisms. But there is a folklore conjecture:

Conjecture 1.6. In the cubic surface case, the symplectic monodromy map

$$\pi_1(\mathcal{M}_{cs}) \rightarrow \pi_0(\mathrm{Symp}(\Sigma_{cs}, \omega_{FS}))$$

is an isomorphism.

See [Sei08, Eva10, LLW22, BLW14, Wu14, Eva11, Smi15, Smi23, Smi22] for related results and context. If true, [Conjecture 1.6](#) would allow our results to apply to a symplectic mapping class group, thus providing an avenue to a Thurstonian theory of a symplectic mapping class group and the “longstanding challenge” of finding analogues of the Nielsen-Thurston classification in this context [Smi15, Section 3.2.2].

Acknowledgments: We warmly thank the following people for helpful conversations: Jason Behrstock, Eliot Bongiovanni, Mark Hagen, Giorgio Mangioni, Jacob Russell, and Alessandro Sisto (hierarchical hyperbolicity); Jonny Evans, Jun Li, and Paul Seidel (symplectic mapping class groups); Michela Artebani, Igor Dolgachev, Shigeyuki Kondō, and Alessandra Sarti (moduli spaces); Matthew Stover (ball quotients); Carlos Serván (mapping class groups of 4-manifolds); Aaron Kim (singularity theory); and Mohith Nagaraju (various topics). AW gratefully acknowledges the support of NSF grant DMS-2142712. Both authors gratefully acknowledge the support of the Simons Foundation.

Tool use: This paper contains no AI-generated text. We used AI only late in the project, after writing the initial draft and lecturing on our results: we used OpenAI’s ChatGPT 5.5 and 5.6 for proofreading, literature search and expository suggestions. It also provided the argument in [Remark D.8](#).

2. BACKGROUND ON GEOMETRY

Here we recall standard results in the minimal level of generality we require. In addition to the many standard books, some of which we cite, there are also accessible course notes available online, such as [Sis, Par, Wri], that may be helpful to anyone learning some of this material for the first time.

2.1. Coarse geometry. Recall the following definitions.

- (1) A geodesic space is a metric space with at least one geodesic joining every pair of points. A geodesic is an isometric embedding of a connected interval into the metric space.
- (2) A geodesic triangle consists of three vertices and a geodesic joining each pair of vertices.
- (3) The triangle is called δ -thin if every point on an edge has distance at most δ to a point on one of the other two edges.
- (4) A geodesic space is called δ -hyperbolic if all its geodesic triangles are δ -thin.
- (5) A geodesic space is called Gromov hyperbolic if it is δ -hyperbolic for some δ .

- (6) A map $f : (X, d_X) \rightarrow (Y, d_Y)$ is called K -coarsely Lipschitz if, for all $x_1, x_2 \in X$,

$$d_Y(f(x_1), f(x_2)) \leq K d_X(x_1, x_2) + K.$$

- (7) It is called a K -quasi-isometric embedding if additionally

$$d_Y(f(x_1), f(x_2)) \geq d_X(x_1, x_2)/K - K.$$

- (8) A subset of a metric space is called K -dense if every point in the metric space has distance at most K to a point in the subspace.
- (9) A map is called a K -quasi-isometry if it is a K -quasi-isometric embedding and the image is K -dense.
- (10) A map is coarsely Lipschitz or a quasi-isometric embedding or a quasi-isometry if there exists K such that the appropriate condition holds.
- (11) A quasi-geodesic in a metric space is a quasi-isometric embedding of a closed connected subinterval of $\mathbb{R}$.
- (12) A path $p : I \rightarrow X$ (where I is a closed connected interval) in a metric space is called a D -unparametrized quasi-geodesic if there is a D -quasigeodesic $q : J \rightarrow X$ and a monotone D -coarsely Lipschitz map $r : I \rightarrow J$ such that $p = q \circ r$.

The term “unparametrized quasi-geodesic”, although quite standard, may seem somewhat misleading, as a parameterization is required for the definition.

For any constants K, δ , there is a constant D such that for any two K -quasi-geodesics with the same endpoints in a δ -hyperbolic space, every point of one has distance at most D to a point of the other [BH99, page 401]. Given two quasi-isometric geodesic metric spaces, one is Gromov hyperbolic if and only if the other is [BH99, page 402].

Remark 2.1. Because of Remark 2.2 below, we will not be able to entirely stay in the realm of geodesic metric spaces. This is an unimportant technical detail that can be safely ignored, but for completeness we now present one of many ways it can be dealt with. (For alternatives, see, for example, [L17, V05].) Given a metric space (X, d) and a number $b > 0$, one can consider the graph $\Gamma_b(X)$ with one vertex v_x for each point x of X , and an edge between two vertices v_x, v_y if $d(x, y) \leq b$. Under mild assumptions the map $x \mapsto v_x$ is a quasi-isometry from X to $\Gamma_b(X)$ [BH99, page 152]. If a metric space (X, d) is not geodesic, we will say that it is Gromov hyperbolic if there is some b such that $x \mapsto v_x$ is a quasi-isometry and $\Gamma_b(X)$ is Gromov hyperbolic.

2.2. Electrification. We define the electrification $(X^!, d^!)$ of a metric space (X, d) , along a family of subsets $\mathcal{Y}$, as follows. The notation “!” is to suggest something that shocks, and $X^!$ may be pronounced “X-zap”.

As a set, $X^!$ is obtained from X by the following process: first, add one point “[Y]” for each $Y \in \mathcal{Y}$, and second, add one copy of the unit interval for each pair (y, Y) with $y \in Y \in \mathcal{Y}$, with one endpoint identified with y and the other with [Y]. We call [Y] the generic point of Y .

Define $d^!$ to be the largest function $X^! \times X^! \rightarrow \mathbb{R}$ that satisfies the following conditions:

- (1) $d^!$ is symmetric and satisfies the triangle inequality;
- (2) the restriction of $d^!$ to $X \times X$ is bounded above by d ; and

- (3) the restriction of $d^!$ to each unit interval is bounded above by the usual metric on the unit interval.

One can check that $d^!$ is a metric, making $(X^!, d^!)$ into a metric space. Furthermore, $d^!(y, [Y]) = 1$ whenever $y \in Y \in \mathcal{Y}$, and $d^!(x, y) = d(x, y)$ whenever $x, y \in X$ and the right side is at most 2.

Remark 2.2. Even under favorable assumptions, it may be that $X^!$ is not a geodesic space. However it is always true that, for any $X^!$ as above of a geodesic metric space X , for any $b > 0$, the map $x \mapsto v_x$ of [Remark 2.1](#) is a quasi-isometry, and for our arguments there is little harm in pretending $X^!$ is geodesic.

The most important fact about electrifications we will use is the following. For a proof, see [\[KR14, Proposition 2.6\]](#) and see the discussion before [\[KR14, Proposition 2.6\]](#) for history and context including [\[Far98, Bow12, MR08, DG, MS21\]](#).

Proposition 2.3. *For all $\delta \geq 0$ there exists a $\delta' \geq 0$ and a $D \geq 0$ such that the electrification of a δ -hyperbolic space along a collection of convex subsets is δ' -hyperbolic.*

Moreover, any geodesic in the original space is a D -unparametrized quasi-geodesic in the electrification.

A convex set is a set containing every geodesic segment with endpoints in the set; but note that the citations show convexity is more than is truly required here.

2.3. Comparison geometry. Consider $\kappa \in \{0, -1\}$. Let M_0 be the Euclidean plane, and M_{-1} be the hyperbolic plane.

If (X, d) is a geodesic metric space, we say that (X, d) is $\text{CAT}(\kappa)$ if every geodesic triangle in X is “at least as thin” as a triangle in M_κ with the same edge lengths [\[BH99, page 158\]](#).

Example 2.4. A complete, simply connected Riemannian manifold of sectional curvature at most κ is $\text{CAT}(\kappa)$ [\[BH99, page 173, 193\]](#).

Remark 2.5. We will make frequent use of the following basic facts.

- (1) Every $\text{CAT}(-1)$ space is also $\text{CAT}(0)$ [\[BH99, page 165\]](#).
- (2) There is a unique geodesic joining every pair of points in a $\text{CAT}(0)$ space [\[BH99, page 160\]](#).
- (3) If (X, d) is $\text{CAT}(0)$ and $C \subset X$ is convex and complete in the induced metric, then the closest point projection

$$p_C : X \rightarrow C$$

is well defined and distance non-increasing [\[BH99, page 176\]](#).

- (4) There exists a $\delta \geq 0$ such that every $\text{CAT}(-1)$ space is δ -hyperbolic [\[BH99, page 399\]](#).

We will also make use of the following, which generalizes the fact that if 3 lines in $\mathbb{R}^n$ meet pairwise orthogonally then they have a point in common.

Lemma 2.6. *Suppose A, B, C are non-empty complete convex subsets of a $\text{CAT}(0)$ space and the closest point projection of any one to any other is their intersection. Then $A \cap B \cap C \neq \emptyset$.*

Proof. $A \cap B \cap C$ is the projection of $A \cap B$ to C , as can be seen using $p_C(A \cap B) \subset p_C(A) \cap p_C(B)$. $\square$

We also have the following strong contraction result for projections.

Lemma 2.7. *For all $\epsilon > 0$ there exists a $\delta > 0$ such that if*

- (1) *C is a closed convex subset of a $\text{CAT}(-1)$ space X , and*
- (2) *x, y are points of X such that the geodesic $[x, y]$ joining them is disjoint from the ϵ -neighborhood of C ,*

then there is a path of length at most δ on the boundary of C joining $p_C(x)$ and $p_C(y)$. In particular, $p_C([x, y])$ has diameter at most δ .

Proof. The second claim is standard. To see the first claim, it suffices to apply the fact that p_C is distance non-increasing to a path obtained from the piecewise geodesic path $[p_C(x), x] \cup [x, y] \cup [y, p_C(y)]$ by replacing the (0, 1, or 2) portions outside the ϵ -neighborhood of C with geodesic segments. $\square$

2.4. Complex hyperbolic geometry. Consider $\mathbb{C}^{n+1}$ with the Hermitian form

$$h(x, y) = -x_0\bar{y}_0 + \sum_{i=1}^n x_i\bar{y}_i.$$

A vector x is called negative if $h(x, x) < 0$. Complex hyperbolic space $\mathbb{CH}^n$ can be defined as the projectivization of the set of negative vectors. This is endowed with a natural Kähler metric called the complex hyperbolic metric, whose group of holomorphic isometries is $PU(n, 1)$. Here $U(n, 1)$ is the group of complex linear maps preserving h , and the linear action of $U(n, 1)$ on $\mathbb{C}^{n+1}$ induces the action of $PU(n, 1)$ on complex hyperbolic space. The complex hyperbolic metric is complete and has sectional curvatures in $[-4, -1]$.

Every negative vector x can be scaled to have $x_0 = 1$, allowing $\mathbb{CH}^n$ to be identified with the ball

$$B = \left\{ z \in \mathbb{C}^n : \sum_{i=1}^n |z_i|^2 < 1 \right\}.$$

A hyperplane in $\mathbb{CH}^n$ is the image under a holomorphic isometry of the locus $\{[x] \in \mathbb{CH}^n : x_n = 0\}$. All hyperplanes are convex and isometric to $\mathbb{CH}^{n-1}$.

More generally, say V is a subspace of $\mathbb{C}^{n+1}$ where the form is positive definite. Then $V^\perp$ defines a convex subspace $\mathbb{CH}(V^\perp)$ isometric to $\mathbb{CH}^{n-\dim_{\mathbb{C}} V}$. All hyperplanes arise in this way with $\dim_{\mathbb{C}} V = 1$.

If $[z]$ is a point in $\mathbb{CH}(V^\perp)$, then the projectivization of $\mathbb{C}z \oplus V$ intersects $\mathbb{CH}^n$ in a sub-ball $\mathbb{CH}(\mathbb{C}z \oplus V)$ that intersects $\mathbb{CH}(V^\perp)$ orthogonally at the point $[z]$. There is a map from the ball to $\mathbb{CH}(V^\perp)$ induced by the orthogonal linear projection map $\mathbb{C}^{n+1} = V \oplus V^\perp \rightarrow V^\perp$. In the above setting it maps $\mathbb{CH}(\mathbb{C}z \oplus V)$ to $[z]$. This map is the closest point projection onto $\mathbb{CH}(V^\perp)$. See for example [Gol99, Section 3.1.5, Lemma 4.3.1].

2.5. Branched covers. We call a collection $\mathcal{H}$ of hyperplanes an orthogonal arrangement if it is locally finite and if any two hyperplanes in $\mathcal{H}$ are either orthogonal or disjoint. We now recall the following result from [All00a, Section 3].

Theorem 2.8. *Let $\mathcal{H}$ be an orthogonal hyperplane arrangement in B , and let $\hat{B}$ be the metric completion of the universal cover of $B - \mathcal{H}$. Then $\hat{B}$ is $\text{CAT}(-1)$.*

There is a natural map $\hat{B} \rightarrow B$, and $\hat{B}$ can be thought of as the universal branched cover of B branched along $\mathcal{H}$. Theorem 2.8 is so fundamental to our

analysis that we will frequently use it without further comment. The following warning will become evident later in the paper, but we state it now for emphasis.

Warning 2.9. $\hat{B}$ is not locally compact. The $\hat{B}$ geodesic joining two points in the universal cover of $B - \mathcal{H}$ does not always stay in that universal cover.

Remark 2.10. Experts in Teichmüller theory may wish to loosely compare $\hat{B}$ to augmented Teichmüller space (also known as the Deligne-Mumford bordification of Teichmüller space).

3. CRITERIA FOR HIERARCHICAL HYPERBOLICITY

In this section we give sufficient criteria for hierarchical hyperbolicity. These criteria have large overlap with the usual definitions and are easy to prove, but are nonetheless convenient for our purposes. The proofs of the criteria are relegated to [Appendix A](#). Non-experts can consult [Wri26, Sis19] for an introduction to hierarchical hyperbolicity if desired, but we do not assume any prior knowledge.

Speaking colloquially, one might say that the skeleton of a hierarchically hyperbolic space is a suitably reasonable metric space equipped with a suitably reasonable collection of maps to (Gromov) hyperbolic spaces. We isolate this core part of the definition of an HHS as follows.

Definition 3.1. We will say that a metric space is quasi-geodesic if there exists q such that every pair of points can be joined by a q -quasigeodesic.

Definition 3.2. Let $(\mathcal{X}, d)$ be a quasi-geodesic space, and let $\mathfrak{S}^b$ be a set. Suppose there exists $\delta, K \geq 0$ such that for each $U \in \mathfrak{S}^b$, we have a δ -hyperbolic space $(\mathcal{C}(U), d_U)$ and a (K, K) -coarsely Lipschitz map

$$\pi_U : \mathcal{X} \rightarrow \mathcal{C}(U).$$

Then we say $(\mathcal{X}, \mathfrak{S}^b, \pi_\bullet)$ is a pre-HHS. Elements of $\mathfrak{S}^b$ are called domains, and $\mathfrak{S}^b$ is called the index set.

We reserve the notation $\mathfrak{S}$ for index sets known to have all of the structure required for hierarchical hyperbolicity, and use $\mathfrak{S}^b$ for index sets otherwise. The notation $\mathfrak{S}^b$ can be pronounced as “Sigma flat”.

From one point of view, the idea of hierarchical hyperbolicity is that $\pi_\bullet$ should be a type of constrained coordinate system for $\mathcal{X}$. The constraints on the coordinates appear as severe restrictions on the maps

$$\pi_U \times \pi_V : \mathcal{X} \rightarrow \mathcal{C}(U) \times \mathcal{C}(V).$$

These restrictions come in one of three forms, and accordingly any pair (U, V) of domains relate to each other in one of three ways. We isolate this combinatorial part of the definition of an HHS as follows.

Definition 3.3. A set $\mathfrak{S}$ is called a valid index set if it is endowed with a symmetric anti-reflexive orthogonality relation $\perp$ and a nesting partial order $\sqsubseteq$ with a unique maximal element, and the following hold:

- (1) **Finite height:** The size of chains for the nesting partial order is bounded.
- (2) **Coherence:** If $V \sqsubseteq W$ and $U \perp W$ then $U \perp V$.
- (3) **Containers:** If $V \sqsubseteq W$ then the set $\mathcal{P}_{V,W}$ of domains orthogonal to V and nested in W is either empty or there is a domain $U \sqsubset W$ such that every domain in $\mathcal{P}_{V,W}$ is nested in U .

If $U, V \in \mathfrak{S}$ are not orthogonal and neither is nested in the other, we say they are asynchronous and write $U \pitchfork V$. (We use “asynchronous” instead of the standard term “transverse”, since the latter is potentially misleading in our examples.)

The fact that $\perp$ is anti-reflexive means that no domain is orthogonal to itself, and together with (2) implies that a nested pair is never orthogonal.

The containers axiom indicates that, when V is nested in W and one restricts attention to domains nested in W , then the set of domains orthogonal to V must be in a sense localized. We will find it useful to weaken this requirement slightly in the following way. Note that we say that a set (or tuple) of domains is orthogonal if every pair of distinct domains in the set (or tuple) is orthogonal.

Definition 3.4. Suppose Definition 3.3 holds except with (3) replaced by

- (3') **Weak containers:** If $V \subseteq W$ then the set $\mathcal{P}_{V,W}$ of domains orthogonal to V and nested in W is either empty or there is an orthogonal set $\{U_i\}$ consisting of domains properly nested in W such that every domain in $\mathcal{P}_{V,W}$ is nested in one of the U_i . Furthermore the size of orthogonal sets of domains is bounded.

Then we say $\mathfrak{S}^b$ is a nearly valid index set.

Example 3.5. If Σ_g is a (2 real dimensional) closed surface of genus g , a subsurface of Σ_g is a closed submanifold U with boundary such that no component of U or $\Sigma_g - U$ is a disc. If $\mathfrak{S}^b$ is the set of connected subsurfaces up to isotopy, then $\mathfrak{S}^b$ is a nearly valid index set. Here U is orthogonal to V if $U \neq V$ and they can be isotoped to be disjoint; and U is nested in V if they are not orthogonal and U can be isotoped to be a subset of V . If $W = \Sigma_g$, then the orthogonal set $\{U_i\}$ required by the weak containers axiom is the set of components of $\Sigma_g - V$ together with (if V is not annular) an annular subsurface for each boundary component of V .

In what follows, we call $\mathcal{A}_U$ the active region for U ; one should imagine that π_U is difficult to change outside of $\mathcal{A}_U$.

Proposition 3.6. *Let $(\mathcal{X}, \mathfrak{S}^b, \pi_\bullet)$ be a pre-HHS with a nearly valid index set. Suppose we fix $B \geq 0$ and a subset $\mathcal{A}_U \subset \mathcal{X}$ for each $U \in \mathfrak{S}^b$, and suppose the following hold.*

- (1) **Enough projections:** *For each $\kappa \geq 0$ there exists $\omega \geq 0$ such that if $d(x, y) \geq \omega$ then $d_U(\pi_U(x), \pi_U(y)) \geq \kappa$ for some U .*
- (2) **Density:** *For all orthogonal tuples $T = (U_1, \dots, U_k)$ of domains,*

$$\pi_{U_1} \times \dots \times \pi_{U_k}(\mathcal{A}_{U_1} \cap \dots \cap \mathcal{A}_{U_k})$$

is B -dense in $\mathcal{C}(U_1) \times \dots \times \mathcal{C}(U_k)$.

- (3) **Bounded projections:** *If $U \pitchfork V$ or $U \subsetneq V$, then*

$$\text{diam } \pi_V(\mathcal{A}_U) \leq B.$$

Furthermore, if $U \subsetneq V$ then $\mathcal{A}_U \cap \mathcal{A}_V \neq \emptyset$.

- (4) **Structured redundancy:** *If $U \pitchfork V$, for any x we have at least one of*

$$d_U(\pi_U(x), \pi_U(\mathcal{A}_V)) \leq B \quad \text{or} \quad d_V(\pi_V(x), \pi_V(\mathcal{A}_U)) \leq B.$$

If $U \subsetneq V$, for any x, y , if there is a geodesic from $\pi_V(x)$ to $\pi_V(y)$ that has distance at least B from $\pi_V(\mathcal{A}_U)$, then $d_U(\pi_U(x), \pi_U(y)) \leq B$.

(5) **Progress localization:** For any $P \geq 0$ there exists $N \geq 0$ and an $R \geq 0$ such that for any V and any x, y with

$$d_V(\pi_V(x), \pi_V(y)) \leq P$$

there is a set of at most N domains $U_i \subsetneq V$ such that if $T \subsetneq V$ is not nested in or orthogonal to one of the U_i then $d_T(\pi_T(x), \pi_T(y)) \leq R$.

Then $\mathcal{X}$ is a hierarchically hyperbolic space.

More precisely, the proof enlarges $\mathfrak{S}^b$ to a larger index set $\mathfrak{S}$ and shows that $(\mathcal{X}, \mathfrak{S}, \pi_\bullet)$ is an HHS. The domains of $\mathfrak{S} - \mathfrak{S}^b$ can be regarded as “dummy domains”; their coordinate functions are constant and their addition can be viewed as a minor technical point. In this case the dummy domains are exactly the orthogonal subsets of $\mathfrak{S}^b$ of size greater than 1, and in the analogy to [Example 3.5](#) they should be compared to disconnected surfaces.

From the coordinate system point of view, speaking very roughly, (1) indicates there are enough coordinates to distinguish points, (2) indicates orthogonal coordinates are independent, (3) indicates that $\mathcal{A}_U$ is collapsed to a point in $\mathcal{C}(V)$, (4) indicates that the π_V coordinate is largely redundant off the corresponding active region $\mathcal{A}_V$, and (5) indicates a sort of local finiteness for the coordinate system.

Even when an assumption of [Proposition 3.6](#) corresponds to an axiom of hierarchical hyperbolicity, all of which have standard names, we have chosen a different name to more clearly indicate when we are referring to an assumption and when we are referring to an axiom.

We also have a variant for groups. Recall that a group action on a metric space is called geometric if the group acts by isometries and the action is properly discontinuous and cocompact.

Proposition 3.7. *Suppose $(\mathcal{X}, \mathfrak{S}^b, \pi_\bullet)$ satisfies the assumptions of [Proposition 3.6](#), $\mathcal{X}$ is proper, and G is a group that acts on $\mathcal{X}$ geometrically. Suppose G also acts on the index set $\mathfrak{S}^b$, via a map $g \mapsto g^\diamond \in \text{Aut}(\mathfrak{S}^b)$, preserving the orthogonality and nesting relations, and there are finitely many orbits of orthogonal tuples. Suppose that for each $U \in \mathfrak{S}^b$, $g \in G$ there is an isometry $g(U) : \mathcal{C}(U) \rightarrow \mathcal{C}(g^\diamond U)$, and*

$$g\mathcal{A}_U = \mathcal{A}_{g^\diamond U} \quad \text{and} \quad g(U) \circ \pi_U = \pi_{g^\diamond U} \circ g.$$

Finally, suppose that for each $U \in \mathfrak{S}^b$ and $g, h \in G$ we have $(hg)(U) = h(g^\diamond U) \circ g(U)$. Then G is a hierarchically hyperbolic group.

4. THE MODEL SPACE AND STATEMENT OF THE GOAL

In this section we will define a proper metric space $\hat{B}^\theta$ on which $\hat{\Gamma}$ acts properly and cocompactly. It will serve as a sort of geometric model for $\hat{\Gamma}$. Our eventual goal in this paper is to apply [Proposition 3.7](#) with $\mathcal{X} = \hat{B}^\theta$ and $G = \hat{\Gamma}$.

Let $s_\theta \gg_{n, \mathcal{H}, \Gamma} 1$ be a very large constant depending on $n, \mathcal{H}, \Gamma$; see [Remark 4.2](#) for details. Having fixed s_θ , for each cusp $c \in \partial B$ of Γ choose a closed horoball T_c , centered at c and small enough that every hyperplane coming within s_θ of T_c passes through c (see for example [\[BH13, page 21\]](#)). The notation “ T ” is for consistency with the next paragraph, where it indicates “tubular”. The T_c are the only horoballs of B that we will consider, so a horoball resp. horosphere (of B) will always mean T_c resp. ∂T_c for some cusp c . We may suppose without loss that the family of horoballs obtained this way is Γ -invariant, and every pair of them have distance at least s_θ from each other.

Remark 4.1. For what comes, it is helpful to note that, since no hyperplane can be contained entirely in a horoball, and since the arrangement is locally finite, there are only finitely many Γ -orbits of hyperplanes.

We may choose a very small constant $0 < r_\theta \ll_{n, \mathcal{H}, \Gamma, s_\theta} 1$ depending on $n, \mathcal{H}, \Gamma, s_\theta$. We define T_H to be the closed r_θ -neighborhood of any hyperplane H of B . (Recall that we have specialized “hyperplane of B ” to mean a member of $\mathcal{H}$.)

Remark 4.2. We require the following of our choices of s_θ and r_θ .

- (1) It is convenient to assume, especially for [Lemma 11.4](#), that any hyperplane that comes within some large distance of one of the horoballs T_c actually enters the T_c . A weaker assumption is used in [Lemma 5.3](#). We also assume that the projection of any horoball to any other horoball has diameter at most 1, for [Lemma 11.4](#), and that the projection of a horoball to any stratum closure not entering the horoball has diameter at most 1, again for [Lemma 11.4](#).
- (2) [Section 5](#) will show that each point of B not in one of the T_c has a small neighborhood in which $\mathcal{H}$ with certain properties. The size of this neighborhood, call it ξ , depends on $n, \mathcal{H}, s_\theta$. We require, especially for [Section 5.2](#), that r_θ times a constant depending on n is smaller than ξ . A weaker assumption is used in [Lemmas 8.11](#) and [8.12](#).
- (3) It would be nice to choose the T_H thin enough to miss each other except near orthogonal intersections. This is not possible, because two hyperplanes may approach each other at a cusp. However we will assume r_θ is small enough that if x and y are on hyperplanes that don’t intersect, and neither x nor y is on a horoball, then the distance from x to y is at least $1000nr_\theta$.

In an attempt to minimize the number of different constants used in the paper, we sometimes pick particular new constants as a function of n and r_θ , even when such a precise choice is not required.

Note that we say two hyperplanes are parallel if they are disjoint but share a point at infinity. Our assumptions imply in particular the following.

- (1) If H, H' are parallel, then $T_H \cap T_{H'} \subseteq T_c$, where the cusp $c \in \partial B^n$ is the common limit point of H and H' in ∂B^n .
- (2) If H, H' are neither parallel nor orthogonal, then $T_H \cap T_{H'} = \emptyset$.

Define B^θ (“ B thick”) as the closure in B of the complement of the union of all T_c and T_H .

Define $\hat{B}^\theta$ as its universal cover, which is also the universal cover of $\Gamma \backslash B^\theta$. Equip $\hat{B}^\theta$ with its lifted intrinsic path metric, and note that since $\Gamma \backslash B^\theta$ is compact we have that $\hat{B}^\theta$ is a proper metric space.

We will show that $\hat{B}^\theta$ is a Hierarchically Hyperbolic Space (HHS) using our criteria for hierarchical hyperbolicity with $\mathcal{X} = \hat{B}^\theta$. Recall that these criteria require a set $\mathfrak{S}^b$ of domains, and for each domain U a Gromov hyperbolic space $\mathcal{C}(U)$ and a coarsely Lipschitz map $\pi_U : \hat{B}^\theta \rightarrow \mathcal{C}(U)$.

We will use four types of domains: “stratum domains”, “angle domains”, “center domains” and “tree domains”. These domains and their associated spaces and projections are defined below.

Remark 4.3. Each of our $\pi_U : \hat{B}^\theta \rightarrow \mathcal{C}(U)$ will be defined on a larger subset $\hat{B}^U$, with

$$\hat{B}^\theta \subset \hat{B}^U \subset \hat{B},$$

giving a map

$$\pi_U : \hat{B}^U \rightarrow \mathcal{C}(U)$$

whose restriction to $\hat{B}^\theta$ is $\pi_U : \hat{B}^\theta \rightarrow \mathcal{C}(U)$. The metric on $\hat{B}^U$ will always be the intrinsic path metric. Details will be provided in [Remarks 5.6, 7.4, 7.6](#) and [8.19](#). In every result about one of the π_U we will indicate which version we mean, but the distinction is largely a technical detail.

5. STRATUM DOMAINS AND THE BRANCHED COVER

In this section we describe the “stratum domains” required for our HHS structure on $\hat{B}^\theta$. The main ingredients in the definition are the universal branched cover $\hat{B}$ of B over $\mathcal{H}$ and various of its subspaces. These are also essential for other parts of the HHS structure.

We will consistently use hats to indicate objects associated to the universal cover of $B^\circ = B - \mathcal{H}$. For example, we define $\hat{B}^\circ$ as this universal cover, $\hat{B}$ as its metric completion, and $\hat{\mathcal{H}} \subseteq \hat{B}$ as the preimage of $\mathcal{H}$. The deck group of $\hat{B}^\circ \rightarrow \Gamma \backslash B^\circ$ has already been named $\hat{\Gamma}$. It is easy to see that B° deformation retracts to B^θ , so we may regard $\hat{B}^\theta$ as the preimage of B^θ in $\hat{B}^\circ \subseteq \hat{B}$. The three spaces $\hat{B}^\theta \subseteq \hat{B}^\circ \subseteq \hat{B}$ lie over the three spaces $B^\theta \subseteq B^\circ \subseteq B$. We think of the natural map $b : \hat{B} \rightarrow B$ as a branched cover.

5.1. Local models. To give a detailed local description, we begin by constructing a special chart around any fixed $x \in B$. Suppose exactly k hyperplanes of B pass through x . Their intersection is a B^{n-k} . The intersection of any $k-1$ of them is a B^{n-k+1} , and we consider the B^1 therein that passes through x and is orthogonal to B^{n-k} . We complete these B^1 ’s to a “set of axes” by choosing $n-k$ additional B^1 ’s inside B^{n-k} , that are mutually orthogonal and pass through x . We now have n many B^1 ’s. The closest point projection to each B^1 , being the projectivization of a linear map, is holomorphic. Nearest-point projection to the product of the B^1 ’s identifies B with an open subset of $(B^1)^n$. (If one takes the B^1 ’s to be the standard coordinate directions through the origin of the ball, then this is the standard inclusion of the ball in $\mathbb{C}^n$ into the product of the n balls in $\mathbb{C}^1$.) We order the factors so that the first k of them are the first k that we specified. We index the hyperplanes through x by $\{i : 1 \leq i \leq k\}$, such that the i th one is the preimage of $(B^1)^{i-1} \times \{x\} \times (B^1)^{n-i}$.

In each B^1 we choose a closed disk centered at x , and define $V \subseteq B$ as the preimage of their product. By choosing the disks small enough, we may suppose that V projects surjectively (hence biholomorphically) to their product. We may also suppose that V is disjoint from gV when $g \in \Gamma - \text{Stab}_\Gamma(x)$, and no hyperplanes intersect V except those containing x . Next we identify each disk biholomorphically with the closed unit disk $D \subseteq \mathbb{C}$, such that x corresponds to the origin. Equipping D with its standard Euclidean metric, and D^n with the product metric, we then have

- (1) V is a closed neighborhood of x , identified with D^n via a map which is biholomorphic and bi-Lipschitz.

- (2) The point x corresponds to $(0, \dots, 0)$.
- (3) Whenever $1 \leq i \leq k$, the intersection of V with the i th hyperplane through x corresponds to $\{(y_1, \dots, y_{i-1}, 0, y_{i+1}, \dots, y_n)\} \subseteq D^n$.
- (4) $V \cap \mathcal{H}$ corresponds to the union of these k many copies of D^{n-1} in D^n .

Remark 5.1. These coordinates have many properties beyond those guaranteed by their bi-Lipschitz and biholomorphic properties. For example, their definition via nearest-point projection implies that the segments in V which are constant in all but one coordinate and are a radial segment in the remaining coordinate are geodesics. These geodesics give the shortest paths to the relevant hyperplane.

Our next step is to define V° as $V - \mathcal{H}$ and work out its preimage in $\hat{B}^\circ$. Writing D° for $D - \{0\}$, V° corresponds to

$$(D^\circ)^k \times D^{n-k} \subseteq D^n.$$

So $\pi_1(V^\circ) \cong \mathbb{Z}^k$. It is easy to see that $\pi_1(V^\circ) \rightarrow \pi_1(B^\circ)$ is injective, by using winding number. Therefore the preimage of V° in $\hat{B}^\circ$ is a disjoint union of copies of the universal cover $\hat{V}^\circ$ of V° . The $\pi_1(B^\circ)$ -stabilizers of these copies are the conjugates of $\pi_1(V^\circ)$. Because $\hat{B}$ was defined as the metric completion of $\hat{B}^\circ$, the inclusion $\hat{V}^\circ \rightarrow \hat{B}^\circ$ extends to a natural map from the metric completion $\hat{V}$ of $\hat{V}^\circ$ to $\hat{B}$. It is easy to see that this is injective, and that the closures of distinct components of $b^{-1}(V^\circ)$ are disjoint. Therefore $b^{-1}(V)$ is a disjoint union of copies of $\hat{V}$. (We work with closed neighborhoods because using open neighborhoods would complicate this statement.)

This reduces the study of the branching above x to understanding $\hat{V}$ and its projection to V . For this we use the bi-Lipschitz property of our identification between V and D^n . Because the metric on $\hat{V}^\circ$ is defined by taking infima of lengths of paths, the bi-Lipschitz identification $V^\circ \cong (D^\circ)^k \times D^{n-k}$ lifts to a bi-Lipschitz identification of $\hat{V}^\circ$ with $(\hat{D}^\circ)^k \times D^{n-k}$. Here $\hat{D}^\circ \subseteq \mathbb{C}$ is the universal cover of D° , namely the closed left half-plane, with $\hat{D}^\circ \rightarrow D^\circ$ being the exponential map. Here $\hat{D}^\circ$ does not have the Euclidean metric, but rather its pullback under this map. The deck group $\mathbb{Z}^k$ consists of the translations by integral multiples of $2\pi i$ in the $\hat{D}^\circ$ factors. The bi-Lipschitz identification $\hat{V}^\circ \cong (\hat{D}^\circ)^k \times D^{n-k}$ extends to an identification of their metric completions. Therefore $\hat{V}$ is bi-Lipschitz to $\hat{D}^k \times D^{n-k}$, where $\hat{D}$ is the metric completion of $\hat{D}^\circ$. It is easy to see that $\hat{D} = \hat{D}^\circ \cup \{-\infty\}$, where the distance between $-\infty$ and any given $z \in \hat{D}^\circ$ is $\exp(\operatorname{Re}(z))$. The map $\hat{D} \rightarrow D$ is still the exponential map, with the convention $\exp(-\infty) = 0$.

This completes our description of $b^{-1}(V) \rightarrow V$, namely: V is bi-Lipschitz to D^n , with x identified with $(0, \dots, 0)$, such that each component of $b^{-1}(V)$ is identified with $\hat{D}^k \times D^{n-k}$. The projection to D^n applies the exponential map to all $\hat{D}$ factors. Each component of $b^{-1}(V)$ contains exactly one preimage of x , corresponding to $(-\infty, \dots, -\infty, 0, \dots, 0)$ with k many $-\infty$'s. The $\pi_1(B^\circ)$ -stabilizers of the preimages of x are the conjugates of $\pi_1(V^\circ) \cong \mathbb{Z}^k$.

5.2. Uniformity. The neighborhoods V cannot be chosen of uniformly large size even if x is assumed to not lie in a horoball, because x can lie arbitrarily close to a hyperplane not containing it.

For each $\hat{V}$ centered at a point $x \in \hat{B}$, let $\hat{V}_0 \subset \hat{V}$ be a small ball centered at x , so that the ball centered at x with radius 100 times bigger is contained in

$\hat{V}$. Assume additionally that the $\hat{V}$ chosen in this way are small enough that the implicit bi-Lipschitz constant in the statement “ $\hat{V}$ is bi-Lipschitz to $\hat{D}^k \times D^{n-k}$ ” is at most 2. This can be achieved by noting that as the size of the neighborhood V decreases, the implicit bi-Lipschitz constant in the statement “ V is bi-Lipschitz to D^n ” goes to 1.

A compactness argument shows that there is a finite collection of Γ orbits of V as above such that each y not in a horoball is contained in a $\hat{V}_0$ of some $\hat{V}$ covering an element of one of those finitely many Γ orbits of V . Let $\xi/100$ be the min of the radii of the $\hat{V}_0$ in this construction. This gives the constant ξ in [Remark 4.2](#).

Remark 5.2. Suppose $x \in \hat{B}$ is not in any horoball. Keeping in mind that [Remark 4.2](#) assumes $r_\theta \ll \xi$, we can see that there is a $x' \in \hat{B}^\theta$ with $d(x, x') \leq 8\sqrt{nr_\theta}$. To see this, first find one of the chosen $\hat{V}_0$ containing x . The associated $\hat{V}$ is identified with $\hat{D}^k \times D^{n-k}$, and proceed as follows.

First, if x is very close to a horoball, push it radially outwards away from the horoball to get a point x'' with $d(x, x'') \leq 4\sqrt{nr_\theta}$ and so that x'' has distance at least $4\sqrt{nr_\theta}$ from every horoball. (Compare to the proof of [Lemma 5.3](#) for details.)

Then we find x' by modifying the first k coordinates of x'' , if required, so that each has distance in $\hat{D}$ more than $2r_\theta$ from the center of $\hat{D}$. Using that the bi-Lipschitz constant is at most 2, this can produce a point x' with $d(x', x'') \leq 4\sqrt{nr_\theta}$ such that x' is not within r_θ of any hyperplane intersecting $\hat{V}$. The triangle inequality gives that x' is not in the interior of any horoball, so $x' \in \hat{B}^\theta$.

5.3. Stratification. The intersections of hyperplanes in B define a stratification of B . Namely, the codimension- k strata are the connected components of the set of points in B that lie in exactly k hyperplanes. This induces a stratification of $\hat{B}$: the strata are the connected components of the preimages of strata of B . In keeping with the use of $^\circ$ to suggest “with hyperplanes removed”, we will usually use notation like S° resp. $\hat{S}^\circ$ for strata in B resp. $\hat{B}$, and S resp. $\hat{S}$ for their closures. Our local model makes clear that the restriction of b to each stratum of $\hat{B}$ is a covering map to a stratum of B . It also shows that the closure $\hat{S}$ of a codimension- k stratum is a component of the fixed point set of a subgroup $\mathbb{Z}^k \subseteq \hat{\Gamma}$. Since $\hat{B}$ is CAT(-1), fixed-point sets are convex, so $\hat{S}$ is the full fixed-point set of this $\mathbb{Z}^k$. Therefore it is convex and CAT(-1).

Just as we reserved the word hyperplane (of B) to mean a component of $\mathcal{H}$, when we speak of hyperplanes (of $\hat{B}$) we mean the closures of codimension 1 strata of $\hat{B}$. These lie over the hyperplanes of B .

5.4. Horoballs. We also need the notion of a horoball of $\hat{B}$. Recalling that we reserved the word horoball (of B) to mean T_c for some cusp c of Γ , a horoball of $\hat{B}$ will mean a component of the preimage of a horoball of B . A horosphere of $\hat{B}$ will mean the boundary of such a horoball. Above each T_c lie infinitely many horoballs of $\hat{B}$; when we have some particular one in mind, we will usually denote it $\hat{T}_{\hat{c}}$, thinking of $\hat{c}$ as a lift of the cusp c . See [Section 7](#) for a detailed analysis of these horoballs; for now we only need their convexity, which we will prove presently in [Lemma 5.4](#).

Lemma 5.3 (Nearest-point projection near a horoball of $\hat{B}$). *Suppose $\hat{T}_{\hat{c}}$ is a horoball of $\hat{B}$ and $\hat{U}_{\hat{c}}$ is the set of points of $\hat{B}$ at distance < 1 from it. Then the*

nearest-point projection map $p_{\hat{T}_c} : \hat{U}_c \rightarrow \hat{T}_c$ is well-defined and enjoys the following property. For each $r > 0$ there is a constant $k < 1$, such that if $\hat{\delta}$ is a path in $\hat{U}_c$ that stays at least r away from $\hat{T}_c$, then $\ell(p_{\hat{T}_c} \circ \hat{\delta}) \leq k\ell(\hat{\delta})$.

Proof. Write T_c for the horoball of B over which $\hat{T}_c$ lies, U_c for the set of points of B at distance < 1 from it, and $p_{T_c} : U_c \rightarrow T_c$ for the nearest-point projection map. Suppose $\hat{x} \in \hat{U}_c$ and set $x = b(\hat{x}) \in U_c$. For our first claim, it is enough to show that there is a unique path from $\hat{x}$ into $\hat{T}_c$ with length $\leq d(x, T_c)$. Existence will be a byproduct of the uniqueness argument. For any such path $\hat{\gamma}$, $b \circ \hat{\gamma}$ must be a path from x to T_c with length $\leq d(x, T_c)$. There is only one: the geodesic $\gamma = \overline{xp_{T_c}(x)}$. By the definition of T_c and Remark 4.2, every hyperplane of B that meets U_c passes through c . Therefore γ lies in the same stratum S° of B as x . So $\hat{\gamma}$ must lie in the same stratum $\hat{S}^\circ$ of $\hat{B}$ as $\hat{x}$, which necessarily lies over S° . Since $\hat{S}^\circ \rightarrow S^\circ$ is a covering space, the lift of γ starting at $\hat{x}$ exists and is unique. This is the unique possibility for $\hat{\gamma}$, proving the well-definedness of $p_{\hat{T}_c} : \hat{U}_c \rightarrow \hat{T}_c$.

For our second claim, we start with the corresponding “unhatted” statement, which holds in complex hyperbolic geometry just as in real hyperbolic geometry. Namely, for every $r > 0$ there is a constant $k < 1$, such that if δ is a path in U_c that stays at least r away from T_c , then $\ell(p_{T_c} \circ \delta) \leq k\ell(\delta)$. From this follows the claim in the special case that $\hat{\delta}$ lies in a metric ball that misses $\hat{\mathcal{H}}$. Concatenating such paths treats the case that $\hat{\delta}$ misses $\hat{\mathcal{H}}$. Then taking limits treats the full claim, since every path $\hat{\delta}$ in $\hat{U}_c$ is a limit of paths in $\hat{U}_c - \hat{\mathcal{H}}$ whose lengths approach $\ell(\hat{\delta})$. $\square$

Lemma 5.4 (Convexity of horoballs of $\hat{B}$). *Every horoball of $\hat{B}$ is convex.*

Proof. Suppose some horoball $\hat{T}_c$ were nonconvex, so some geodesic $\hat{\gamma}$ in $\hat{B}$ has endpoints in $\hat{T}_c$ but does not lie entirely within $\hat{T}_c$. Recall that geodesics in a CAT(−1) space depend continuously on their endpoints [BH99, page 160]. Moving one endpoint toward the other, inside $\hat{T}_c$, homotopes $\hat{\gamma}$ to a constant geodesic. So some intermediate geodesic exits $\hat{T}_c$ and then returns, but also remains at distance less than 1 from it. Using it in place of $\hat{\gamma}$, we may suppose without loss that $\hat{\gamma}$ has these properties.

Now we apply Lemma 5.3: $\hat{\gamma}$ lies in $\hat{U}_c$, and the fact that it leaves $\hat{T}_c$ implies $\ell(p_{\hat{T}_c} \circ \hat{\gamma}) < \ell(\hat{\gamma})$. Since $p_{\hat{T}_c} \circ \hat{\gamma}$ has the same endpoints as $\hat{\gamma}$, the latter cannot be a geodesic, contrary to hypothesis. $\square$

Remark 5.5. An alternative but slightly less elementary way to prove Lemma 5.4 is to prove $\hat{T}_c$ really is a horoball in the usual sense of [BH99, page 267], by lifting the flow towards the cusp, and then to use that horoballs in CAT(0) spaces are convex [BH99, page 271].

5.5. Stratum domains. Now we can define the stratum domains. They are indexed by the strata $\hat{S}^\circ$ in $\hat{B}$ for all $\hat{S}^\circ$ of positive dimension; we choose not to define stratum domains for 0-dimensional strata. If $\hat{S}$ is the closure of a d -dimensional stratum $\hat{S}^\circ$, then the Gromov hyperbolic space $\mathcal{C}(\hat{S}^\circ)$ associated to $\hat{S}^\circ$ is defined as the electrification $\hat{S}^!$ of $\hat{S}$ along

- (1) the closures of the $(d - 1)$ -dimensional strata that lie in $\hat{S}$, and
- (2) the nonempty intersections of $\hat{S}$ with horoballs of $\hat{B}$.

[Lemma 5.7](#) below shows that $\hat{S}^!$ is Gromov hyperbolic.

An HHS structure also requires a coarsely Lipschitz projection map $\pi_{\hat{S}} : \hat{B}^\theta \rightarrow \mathcal{C}\hat{S}$, which we take to be the composition

$$\pi_{\hat{S}} : \hat{B}^\theta \subseteq \hat{B} \rightarrow \hat{S} \subseteq \hat{S}^!.$$

The map $\hat{B} \rightarrow \hat{S}$ is closest-point projection, which makes sense because $\hat{B}$ is CAT(−1) and $\hat{S}$ is convex. This map and the two inclusions are distance non-increasing, so $\pi_{\hat{S}}$ is certainly coarsely Lipschitz.

Remark 5.6. Continuing [Remark 4.3](#), we note that the true domain of this map is $\hat{B}^{\hat{S}} = \hat{B}$.

Lemma 5.7 (Gromov hyperbolicity of stratum domains). *If $\hat{S}$ is any stratum closure in $\hat{B}$, then $\hat{S}^!$ is Gromov hyperbolic.*

Proof. We observed above that $\hat{S}$ is CAT(−1), hence Gromov hyperbolic. By [Proposition 2.3](#), it is enough to show that each of the subsets of $\hat{S}$ along which we electrify to get $\hat{S}^!$ is convex. First, all stratum closures are convex in $\hat{B}$, so those that lie in $\hat{S}$ are convex there. Second, by [Lemma 5.4](#), all horoballs of $\hat{B}$ are convex in $\hat{B}$. So the horoballs that meet $\hat{S}$ do so in convex subsets. $\square$

Remark 5.8. If the union of the hyperplanes in $\hat{B}$ is connected and either $n \geq 3$ or Γ is cocompact, one can show that the “top-level” hyperbolic space $\mathcal{C}(\hat{B}^\circ)$ is quasi-isometric to the intersection graph of hyperplanes of $\hat{B}$. (If $n = 2$ and Γ is non-uniform, one should add vertices corresponding to the cusps.)

As recalled and reinterpreted in [[ACT02a](#), Section 7], Libgober [[Lib78](#)] showed that in the case of the moduli space of cubic surfaces $\hat{\Gamma}$ is generated by a finite set of elements that each stabilize a hyperplane in $\hat{B}$, such that the union of the hyperplanes thus obtained is connected. Keeping in mind there is only one orbit of hyperplanes [[ACT02a](#), Theorem 7.21], this shows that the union of the hyperplanes of $\hat{B}$ is connected in this case.

Thus, a more combinatorial and geometric top-level hyperbolic space is available in the cubic surfaces case, which is more closely analogous to classical curve graphs (or curve complexes).

Remark 5.9. If $\hat{S}$ enters a horoball $\hat{T}_\epsilon$, then the closest point projection of $\hat{S}$ to $\hat{T}_\epsilon$ is $\hat{S} \cap \hat{T}_\epsilon$, since $\hat{S}$ is the intersection of fixed point sets of isometries preserving $\hat{T}_\epsilon$.

Additionally, the projection of $\hat{T}_\epsilon$ to $\hat{S}$ is $\hat{S} \cap \hat{T}_\epsilon$, because, as just established, projection to $\hat{T}_\epsilon$ preserves $\hat{S}$ and because this projection is distance non-increasing.

Remark 5.10. Additionally using that commuting elements preserve each others fixed point sets, one can see that if $\hat{S}$ and $\hat{S}'$ are intersecting stratum closures, then $p_{\hat{S}}(\hat{S}') = \hat{S} \cap \hat{S}'$.

6. SWADDLING AND CURVATURE

In this section, which is largely self-contained, we develop a sort of theory of $\mathbb{R}$ -bundles, first over arbitrary spaces and then in a more geometric setting.

6.1. Swaddling. We start by introducing swaddling, which coarsely identifies all the different fibers, and give a simple condition under which the result is a quasi-line (a space quasi-isometric to $\mathbb{R}$).

Definition 6.1. Consider a map $E \rightarrow X$ of sets, and assume we specify an action of $\mathbb{R}$ on E which preserves fibers and acts simply transitively on each fiber. (So, this is a principal $\mathbb{R}$ -bundle but without topology.) Assume that for each pair $(x, y) \in X^2$ we have an $\mathbb{R}$ -equivariant map

$$m_{x,y} : E_x \rightarrow E_y$$

from the fiber E_x over x to the fiber E_y over y . We assume $m_{y,x} = m_{x,y}^{-1}$ for all $x, y \in X$, which implies that $m_{x,x}$ is the identity for all $x \in X$. We call this data $(E \rightarrow X, m_{\bullet,\bullet})$ swaddling data.

Definition 6.2. Given a choice of swaddling data, we define the swaddled space E^{sw} as follows. For each pair $x, y \in X$ of distinct points, consider a strip $S_{x,y} = [0, 1] \times \mathbb{R}$, equipped with the usual action of $\mathbb{R}$ on the second coordinate by translation. We fix a $\mathbb{R}$ -equivariant identification between $\{0\} \times \mathbb{R}$ and E_x , and also between $\{1\} \times \mathbb{R}$ and E_y , such that via these identifications $m_{x,y}((0, t)) = (1, t)$. We define E^{sw} to be the disjoint union of all the fibers E_x and all the strips $S_{x,y}$ with the specified identifications between fibers and edges of strips.

We equip E^{sw} with the obvious path metric, which can be defined formally as follows: Each point of E^{sw} has a neighborhood which is either an open interval cross $\mathbb{R}$ or the join of half-open intervals cross $\mathbb{R}$. We equip these neighborhoods with the usual product metric, and note that this gives an open cover of E^{sw} with a choice of metric on each open set in the cover such that these choices agree on overlaps. We equip E^{sw} with the associated path metric, as in [BH99, page 32]. (The details of this metric will not be important for us.)

There is a natural inclusion map $E \rightarrow E^{\text{sw}}$. If x, y are distinct points of X , the distance between their fibers in E^{sw} is 1. This is the sense in which we think of swaddling as “weakly” identifying fibers with each other. This compression of E , in all directions except the fiber direction (“head-to-toe”), is the source of the word *swaddled*.

Definition 6.3. For $M \geq 0$, say the swaddling data $(E \rightarrow X, m_{\bullet,\bullet})$ is M -almost trivial if for all $x, y, z \in X$, the $\mathbb{R}$ -equivariant map

$$m_{z,x} \circ m_{y,z} \circ m_{x,y} : E_x \rightarrow E_x,$$

which must be translation by some real number, is a translation of size at most M .

Lemma 6.4. Suppose swaddling data $(E \rightarrow X, m_{\bullet,\bullet})$ is M -almost trivial for some $M \geq 0$. Then the inclusion of any fiber into the swaddled space E^{sw} is a quasi-isometry.

Note that each fiber, being canonically identified with $\mathbb{R}$ up to translations, has the usual metric on $\mathbb{R}$. In particular, we get that E^{sw} is a quasi-line (a space quasi-isometric to $\mathbb{R}$).

Proof. Let us consider the fiber of $x \in X$. The inclusion of the fiber of x is distance non-increasing, and every point of E^{sw} lies within distance $\frac{3}{2}$ of this fiber. So it suffices to build a quasi-inverse r to this inclusion (see for example [L17, Proposition 5.1.10]).

Consider the disjoint union of the fibers of $E \rightarrow X$, which is a $\frac{1}{2}$ -dense subset of E^{sw} . We define a map r from this subset to the fiber of x as follows: For any point e of E , lying over a point $y \in X$, we define $r(e) = m_{y,x}(e)$.

To see that r is coarsely Lipschitz, it suffices to consider two points of E that differ by a map $m_{y,z}$. These have distance at most 1 in E^{sw} , and we must show their images have uniformly bounded distance in the fiber of x . However the assumption exactly gives that their images have distance at most M . $\square$

Above E and X are just sets, but in many cases they will come with metrics, and we note the following metric compatibility result, whose proof is left as an exercise.

Lemma 6.5. *Suppose E is equipped with a path metric, and there exists $\epsilon > 0$ and $C > 0$ such that for all $e, f \in E$ over points $x, y \in X$ we have that $d(e, f) \leq \epsilon$ implies that the unique t such that*

$$m_{x,y}(e) = t \cdot f$$

has $|t| \leq C$. Then the inclusion $E \rightarrow E^{\text{sw}}$ is coarsely Lipschitz.

6.2. Other points of view (not required). Given a central extension

$$1 \rightarrow Z \rightarrow G \rightarrow Q \rightarrow 1$$

any choice of section $s : Q \rightarrow G$ gives rise to a Z -valued 2-cocycle c on Q defined by

$$c(q_1, q_2) = s(q_1 q_2)^{-1} s(q_1) s(q_2).$$

The class $[c] \in H^2(Q, Z)$ is the extension class of the central extension, and gives rise to the well-known bijective correspondence between central Z -extensions up to equivalence and $H^2(Q, Z)$. Especially when $Z = \mathbb{Z}$, the extension class is also known as the Euler class.

Assume Z is equipped with a suitable metric, for example the word metric on a finitely generated group or the usual metric on $\mathbb{R}$. A 2-cocycle c is called bounded if $c(Q, Q) \subset Z$ is bounded, and an element of $H^2(Q, Z)$ is called bounded if it can be represented by a bounded 2-cocycle. A central extension is called bounded if its extension class is bounded.

If $Z = \mathbb{R}$, then we can take $E = G$ and $X = Q$ in the discussion of swaddling. Since $Z = \mathbb{R}$, we can arrange for $s(q^{-1}) = s(q)^{-1}$. Under these assumptions, the section s determines swaddling data as follows: For every $x, y \in X^2$, there is a unique $q \in Q$ such that $y = qx$, and we set $m_{x,y}$ to be left multiplication by $s(q)$ restricted to E_x .

The map

$$m_{q_2 q_1 \cdot x_0, x_0} \circ m_{q_1 \cdot x_0, q_2 q_1 \cdot x_0} \circ m_{x_0, q_1 \cdot x_0} : E_{x_0} \rightarrow E_{x_0}$$

is translation by $c(q_2, q_1) \in Z$, and c is bounded if and only if the swaddling data is almost trivial.

Under appropriate assumptions one can produce central extensions from the swaddling setup. Even though these assumptions do not always hold, [Lemma 6.4](#) should be seen as closely related to Gersten's result [[Ger92](#), Theorem 3.1] that bounded central $\mathbb{Z}$ -extensions are quasi-isometric to direct products $\mathbb{Z} \times Q$. See also [[FS23](#), [FFMS26](#)] for more results related to bounded central extensions.

Quasi-lines can also be built from quasi-morphisms [[ABO19](#), Lemma 4.15], and, following [[HRSS25](#)], this has become a valuable tool in the theory of hierarchical

hyperbolicity [HMS24, DDLS24, FFMS26, Tao25]. Conversely quasi-morphisms can be defined using appropriate actions on quasi-lines [Man08, Section 4.1].

[FFMS26, Proposition 2.9] characterizes boundedness of central extensions in terms of quasi-homomorphisms. We can write every element of G uniquely as $z \cdot s(q)$ for some $z \in Z$ and $q \in Q$. Consider the map

$$\phi : G \rightarrow Z, \quad \phi(z \cdot s(q)) = z,$$

and compute, for $g_i = z_i s(q_i)$,

$$\phi(g_1 g_2) - \phi(g_1) - \phi(g_2) = c(q_1, q_2).$$

Thus c is bounded if and only if ϕ is a quasi-homomorphism.

Remark 6.6. All quasi-lines are by definition quasi-isometric to each other. In our case we think of results like Lemma 6.4 as inseparable from the definition of the space that turns out to be a quasi-line, and keep in mind that this definition provides tools to build and study maps to the quasi-lines that are produced. In particular, in our applications the space E is often quite large – sometimes not even admitting a natural cobounded group action – and the fact that there is a natural map from the large space E to the quasi-line E^{sw} is helpful for our analysis.

6.3. L -bounded curvature. We now study swaddling data which can be thought of as coming from parallel transport along geodesics in a negatively curved base.

Definition 6.7. We say swaddling data $(E \rightarrow X, m_{\bullet, \bullet})$ is geometric if X is equipped with a $\text{CAT}(-1)$ metric, and, whenever $x, y, z \in X$ with y on the geodesic from x to z , we have $m_{x,z} = m_{y,z} \circ m_{x,y}$.

Definition 6.8. Given a triangle with vertices x, y, z in a $\text{CAT}(-1)$ space X , and swaddling data $(E \rightarrow X, m_{\bullet, \bullet})$, we say the monodromy of the triangle with vertices x, y, z is the amount by which

$$m_{z,x} \circ m_{y,z} \circ m_{x,y} : E_x \rightarrow E_x$$

translates. This is a real number that, up to sign, depends only on $\{x, y, z\}$.

Definition 6.9. Given swaddling data $(E \rightarrow X, m_{\bullet, \bullet})$ with X a $\text{CAT}(-1)$ space, and a non-negative number L , we say that a triangle in X is L -bounded if the monodromy around it has magnitude bounded above by L times the area of the comparison triangle in $\mathbb{H}^2$.

The swaddling data itself is called L -bounded if the same is true for all triangles.

Definition 6.10. Given a triangle T in a $\text{CAT}(-1)$ space with vertices x, y, z and a point m on the edge of T opposite y , the splitting of T determined by m is the two triangles with vertices y, m, x and y, m, z respectively.

A subdivision of T is a collection of triangles obtained starting from T by iteratively replacing a triangle T' in the collection with the two triangles obtained by splitting T' .

In the next subsection we will prove the following.

Proposition 6.11. *Suppose that $(E \rightarrow X, m_{\bullet, \bullet})$ is geometric swaddling data, and fix $L \geq 0$. If every point in X has a neighborhood such that every triangle in this neighborhood admits a subdivision into L -bounded triangles, then $(E \rightarrow X, m_{\bullet, \bullet})$ is L -bounded.*

6.4. Subdivisions of triangles in $CAT(-1)$ spaces. The purpose of this subsection is to prove [Proposition 6.11](#).

The advantage of the definition of subdivision we use is the following.

Lemma 6.12 (Super-additivity of areas). *Given a triangle T in a $CAT(-1)$ space and a subdivision of T into triangles $T_1, \dots, T_k$, the area of the comparison triangle for T is at least the sum of the areas of the comparison triangles for the T_i .*

Here is some intuition for this statement, speaking in very imprecise terms: The area of a triangle gets smaller when curvature is more negative. The $CAT(-1)$ space can have more negative curvature than the hyperbolic plane. The subdivision can see more of the negative curvature beyond what is guaranteed by $CAT(-1)$, and so allows for a smaller area.

Proof. Consider two triangles in $\mathbb{H}^2$, with a common edge and disjoint interiors. Write a, w for the vertices at the end of their common edge, and suppose the sum of the angles at w is at least π . Write v, x for the remaining vertices. (See [Figure 1](#).) Then the sum of the areas of these triangles is bounded above by the area of a triangle in $\mathbb{H}^2$ whose edge lengths are $d(a, v)$, $d(v, w) + d(w, x)$ and $d(x, a)$. This is an immediate corollary of Alexandrov's Lemma [[BH99](#), page 25], as was noted in [[CHRT25](#), Lemma 4.5].




FIGURE 1. A schematic of the proof of [Lemma 6.12](#).

Keeping in mind the triangle inequality for Alexandrov angles [[BH99](#), page 10] and [[BH99](#), page 161, Proposition 1.7(4)] this shows the result when $k = 2$, and the result follows by induction on k . $\square$

The definition of subdivision is flexible enough to allow the following.

Lemma 6.13 (Alexandrov patchwork). *Let X be a $CAT(-1)$ metric space, and let $\mathcal{U}$ be an open cover of X . Then every triangle T in X has a subdivision into triangles each of which is contained in an element of $\mathcal{U}$.*

Sketch of proof of [Lemma 6.13](#). Suppose Δ is a geodesic triangle with vertices a, v and x . Consider the surface S swept out by the $\overline{aw}$ as w moves along $\overline{vx}$.

For each point s of S pick a $\delta_s > 0$ such that the ball of radius $2\delta_s$ about s is contained in an element of $\mathcal{U}$. By compactness, the covering of S by the open balls of radius δ_s about all the $s \in S$ has a finite subcover. It follows that there is some $\lambda > 0$ such that every point s of S has that the ball of radius λ about s is contained in an element of $\mathcal{U}$.

So every geodesic triangle in X whose three vertices lie within λ of some point of S is contained in a ball which is contained in an element of $\mathcal{U}$. (Here we use that balls are convex.) The Alexandrov patchwork argument, used for example in [BH99, page 199-200] and [AKP24, pages 102-104], subdivides Δ into such triangles. $\square$

Proposition 6.14. *Fix geometric swaddling data. If a triangle T can be subdivided into L -bounded triangles, then T is L -bounded.*

Proof. We call the triangles in the subdivision “small triangles”. Lemma 6.12 gives that the area of the comparison triangle for T is at least the sum of the areas of the comparison triangles for small triangles. The monodromy of T is the sum of the monodromies of the small triangles (after making appropriate orientation conventions). The proposition follows immediately. $\square$

Proof of Proposition 6.11. This follows immediately from Lemma 6.13 and Proposition 6.14. $\square$

6.5. Pullbacks. Our goal is now to show that frequently L -bounded curvature is stable under pullbacks.

Definition 6.15. Suppose $\hat{X}$ and X are $\text{CAT}(-1)$ spaces. Then we will say a map $b : \hat{X} \rightarrow X$ is piecewise geodesic if every geodesic segment in $\hat{X}$ can be broken into finitely many sub-geodesic segments that each map isometrically to their image in X .

Note that any piecewise geodesic map is 1-Lipschitz.

Lemma 6.16. *Let $b : \hat{X} \rightarrow X$ be a piecewise geodesic map of $\text{CAT}(-1)$ spaces. Suppose $(E \rightarrow X, m_{\bullet, \bullet})$ is geometric swaddling data. Then there is a unique family of maps $\hat{m}_{\bullet, \bullet}$ such that $(b^*(E) \rightarrow \hat{X}, \hat{m}_{\bullet, \bullet})$ is geometric swaddling data for the pullback, and such that if $x, y \in \hat{X}$ is such that the geodesic from x to y maps isometrically to its image in X , then $\hat{m}_{x, y} = m_{b(x), b(y)}$.*

By definition, the fiber of the pullback bundle $b^*(E)$ over a point $x \in \hat{X}$ is naturally identified with the fiber of E over $b(x)$, and it is via such identifications that the equality $\hat{m}_{x, y} = m_{b(x), b(y)}$ should be interpreted.

Proof. For any $x, y \in \hat{X}$, there exists a sequence $x = x_0, x_1, \dots, x_k = y$ of points on the geodesic $[x, y]$ such that each geodesic $[x_i, x_{i+1}]$ maps isometrically to a geodesic in X . We define

$$\hat{m}_{x, y} = m_{b(x_{k-1}), b(x_k)} \circ \dots \circ m_{b(x_0), b(x_1)}.$$

Because the original swaddling data $(E \rightarrow X, m_{\bullet, \bullet})$ is geometric, this definition is stable under adding extra points to the sequence. Taking common subdivisions shows that $\hat{m}_{x, y}$ is well-defined. Uniqueness is clear. $\square$

Definition 6.17. We define the swaddling data $(b^*(E) \rightarrow \hat{X}, \hat{m}_{\bullet, \bullet})$ produced by Lemma 6.16 to be the pullback swaddling data.

Definition 6.18. Suppose $\hat{X}$ and X are $\text{CAT}(-1)$ spaces. Then we will say a map $b : \hat{X} \rightarrow X$ is piecewise triangular if every triangle in $\hat{X}$ can be subdivided, in the sense of Definition 6.10, into (smaller) triangles such that the restriction of b to each edge of one of these (smaller) triangles is an isometric embedding.

Note that piecewise triangular implies piecewise geodesic, since any geodesic segment is the edge of a triangle. We now have our main goal for this section.

Corollary 6.19. *Suppose that $(E \rightarrow X, m_{\bullet, \bullet})$ is geometric swaddling data that is L -bounded, and suppose that $b : \hat{X} \rightarrow X$ is piecewise triangular. Then the pullback swaddling data $(b^*(E) \rightarrow \hat{X}, \hat{m}_{\bullet, \bullet})$ is L -bounded. In particular, it is almost trivial, and hence the swaddled space $b^*(E)^{\text{sw}}$ is a quasi-line.*

Proof. It follows from [Proposition 6.14](#) and the definitions that $(b^*(E) \rightarrow \hat{X}, \hat{m}_{\bullet, \bullet})$ is L -bounded, and it is immediate from the definitions that L -bounded implies almost trivial. The final claim, that the swaddled space is a quasi-line, is provided by [Lemma 6.4](#). $\square$

7. CUSP DOMAINS

Our goal in this section is to continue to fulfill our promise made in [Section 4](#), by defining more of the domains in the HHS on $\hat{B}^\theta$, and the projections from $\hat{B}^\theta$ to them. The domains defined in this section, called “center domains” and “tree domains”, arise from the horoballs of $\hat{B}$. Each horoball will give rise to $n - 1$ tree domains and 1 center domain.

In [Section 4](#) we introduced the horoballs $T_c \subseteq B$ centered at the cusps c of Γ . As usual, we add a superscript $^\circ$ to indicate “minus the hyperplanes”, i.e. $T_c^\circ = T_c \cap B^\circ$. By ∂T_c° we mean $(\partial T_c)^\circ = \partial T_c \cap B^\circ$, not $\partial(T_c^\circ)$.

We continue to write $\hat{T}_{\hat{c}}$ for a horoball of $\hat{B}$ lying over T_c , where $\hat{c}$ indicates one of the many cusps of $\hat{B}$ that lie over c . We attach a superscript $^\circ$ as before, i.e. $\hat{T}_{\hat{c}}^\circ$ and $\partial \hat{T}_{\hat{c}}^\circ$ are the intersections of $\hat{T}_{\hat{c}}$ and $\partial \hat{T}_{\hat{c}}$ with $\hat{B}^\circ$.

We will show that $\partial \hat{T}_{\hat{c}}$ is a type of $\mathbb{R}$ -bundle over a product of $n - 1$ Gromov hyperbolic spaces $\hat{\mathbb{C}}_i$. The $\hat{\mathbb{C}}_i$ are quasi-isometric to electrifications of trees, hence the term “tree domain”. A swaddling of $\partial \hat{T}_{\hat{c}}$ is related to the center of the Heisenberg group, hence the term “center domain”. Formal definitions require some background about the geometry of a complex hyperbolic horosphere.

7.1. The Heisenberg group and horoballs. Here we follow the discussion in [\[All98, Section 3\]](#), but see also [\[Gol99, Chapter 4\]](#) and [\[Par, Section 4\]](#) for details.

Let H^{2n-1} denote the $(2n - 1)$ -dimensional Heisenberg group, and let $Z = Z(H^{2n-1})$ denote its center. This center is isomorphic to $\mathbb{R}$, and the quotient H^{2n-1}/Z is isomorphic to $\mathbb{R}^{2n-2}$ as a Lie group.

Start by supposing c lies in ∂B^n and T_c is a horoball with center c . Consider the bounding horosphere ∂T_c . The $\text{PU}(n, 1)$ -stabilizer of ∂T_c is the semidirect product of the $(2n - 1)$ -dimensional Heisenberg group H^{2n-1} by the unitary group $U(n - 1)$. The factor H^{2n-1} acts simply transitively on ∂T_c , with the $U(n - 1)$ factor being any chosen point stabilizer. Under the action of the center of H^{2n-1} ,

$$Z = Z(H^{2n-1}) \cong \mathbb{R},$$

∂T_c becomes a principal $\mathbb{R}$ -bundle. The base of this bundle is the set of Z -orbits, which can be identified with the set of complex lines in $\mathbb{P}(\mathbb{C}^{n,1})$ that pass through c and meet B^n . Another way to think of the base is as a torsor for $H^{2n-1}/Z \cong \mathbb{C}^{n-1}$. We write $\mathbb{C}^{n-1}$ and not $\mathbb{R}^{2n-2}$, because the induced action of $U(n - 1)$ identifies the quotient group with complex Euclidean space $\mathbb{C}^{n-1}$ (up to a scale factor). The

group structure of the central extension

$$1 \rightarrow Z \rightarrow H^{2n-1} \rightarrow \mathbb{C}^{n-1} \rightarrow 1$$

is determined by the commutator map $H^{2n-1} \times H^{2n-1} \rightarrow Z$. This carries the same information as the antisymmetric bilinear form $\mathbb{C}^{n-1} \times \mathbb{C}^{n-1} \rightarrow Z$ that it induces. By $U(n-1)$ symmetry, this bilinear form must be the imaginary part of the Hermitian inner product on $\mathbb{C}^{n-1}$, up to a constant factor.

Having described H^{2n-1} as a group, next we examine it as a Riemannian manifold. We already know that each fiber is isometric to $\mathbb{R}$, and the base to $\mathbb{C}^{n-1}$. Additional information is carried by the field of hyperplanes orthogonal to the Z -orbits, in the tangent spaces of points of ∂T_c . Usually one imagines the fiber direction as vertical, so these hyperplanes and the vectors they contain are called *horizontal*.

A smooth path in H^{2n-1} is called horizontal if its tangent vectors are. The horizontal hyperplanes form an Ehresmann connection. In particular, given any smooth path γ in the base, and any preimage in H^{2n-1} of its starting point, there is a unique horizontal lift $\tilde{\gamma}$ of γ starting there. If γ is a loop, then the endpoints of $\tilde{\gamma}$ lie in a single fiber. So there is a unique element of Z that sends the starting point to the endpoint, called the *monodromy* around γ .

The monodromy can be expressed in terms of the curvature of the connection. This is a 2-form on the base, taking values in the Lie algebra of the structure group, and describes “the monodromy around infinitesimal rectangles”. In our setting, the structure group is $\mathbb{R}$, so the curvature is an ordinary 2-form on (the real manifold underlying) $\mathbb{C}^{n-1}$.

To find the monodromy around a curve γ , integrate the curvature form over any disk with boundary γ . By $U(n-1)$ symmetry, the only possibility for the curvature form is the imaginary part of the Hermitian form, up to a scale factor. This makes the monodromy easy to understand: choose any orthonormal basis for $\mathbb{C}^{n-1}$ and write $\gamma = (\gamma_1, \dots, \gamma_{n-1})$ with respect to it. Then, up to a global factor, the monodromy around γ is the sum of the signed areas enclosed by the γ_i .

For later use we record the following. In this lemma T_c may be any horoball of B (there is only one up to $PU(n, 1)$).

Lemma 7.1. *There is a constant C depending only on n such that the following holds. Let $\alpha : [0, 1] \rightarrow \partial T_c$ be a C^1 path. Let v_0 be a unit horizontal tangent vector at $\alpha(0)$, and let $\bar{v}_0$ be its image in $\mathbb{C}^{n-1}$. Let v_1 be the parallel translate of v_0 to $\alpha(1)$ along α and let v'_1 be the horizontal lift to $\alpha(1)$ of $\bar{v}_0$. Then the distance between v_1 and v'_1 is at most C times the length of α .*

Here, when $\bar{v}_0$ is viewed as in $\mathbb{C}^{n-1}$, we are identifying the $\mathbb{C}^{n-1}$ as being simultaneously the tangent space to every point of the $\mathbb{C}^{n-1}$ that is ∂T_c mod center.

Proof. There is a vector field V_0 on ∂T_c whose value at a point is the horizontal lift of $\bar{v}_0$.

The action of H^{2n-1} on ∂T_c induces an action by translations on ∂T_c mod center. Since the H^{2n-1} action is transitive, this shows that V_0 is invariant under H^{2n-1} . Since the Levi-Civita connection ∇ is invariant under isometries, we get that ∇V_0 is invariant under H^{2n-1} . In particular, the operator norm of ∇V_0 is constant on ∂T_c .

This gives the result because the difference between v_1 and v'_1 can be computed using the integral of ∇V_0 applied to $\alpha'(t)$. $\square$

7.2. The hyperplane arrangement. We will also need to understand the hyperplanes of B that contain c . Each is the orthogonal complement of a vector $r \in \mathbb{C}^{n,1}$ that satisfies $r^2 > 0$ and $r \perp c$, and we note the following, where “hyperplane” of $\mathbb{C}^{n-1}$ means an affine complex codimension 1 subspace of $\mathbb{C}^{n-1}$, and where we temporarily suspend our standing convention that a hyperplane of B always means a hyperplane in $\mathcal{H}$.

Lemma 7.2. *The map $H \mapsto (H \cap \partial T_c)/Z$ induces a bijection between hyperplanes of B containing c and hyperplanes of $\partial T_c/Z \simeq \mathbb{C}^{n-1}$. Furthermore,*

- (1) *two such hyperplanes H_1 and H_2 are orthogonal in B if and only if the associated hyperplanes in $\mathbb{C}^{n-1}$ are orthogonal, and*
- (2) *if v is a normal vector to the hyperplane associated to H , then any horizontal lift of v is orthogonal to H .*

Since this is standard, we offer only a partial sketch. Say $r \in \mathbb{C}^{n,1}$ is a positive vector. Note that $\partial T_c \cap r^\perp$ is a horosphere in $r^\perp$. The $\mathrm{PU}(n, 1)$ -stabilizer of r is a copy of $\mathrm{PU}(n-1, 1)$, and c 's stabilizer therein is a copy of $H^{2n-3} \rtimes U(n-2)$. Because the center of H^{2n-3} is also the center of H^{2n-1} here, it follows that each hyperplane through c meets ∂T_c in the union of fibers over a hyperplane in $\mathbb{C}^{n-1}$. We leave it to the reader to check orthogonality either using symmetry or by working in coordinates. The final claim follows because the whole horosphere is orthogonal to any geodesic to c , and the geodesic direction and the center direction are the only directions not seen in $\mathbb{C}^{n-1}$.

Now we specialize to our situation. Namely, c is a cusp of Γ , and T_c is the horoball, centered there, that we defined in Section 4. Consider the components of $\mathcal{H}$ that meet c . By the previous paragraph, they correspond to a union of hyperplanes in $\mathbb{C}^{n-1}$. Since $\mathcal{H}$ is an orthogonal arrangement, so is this arrangement in $\mathbb{C}^{n-1}$.

By Assumption (D), c is a limit of some 1-dimensional stratum of B . This stratum is the intersection of $n-1$ hyperplanes of B , all containing c . By $U(n-1)$ symmetry, we may suppose that the corresponding hyperplanes are parallel to the coordinate hyperplanes in $\mathbb{C}^{n-1}$.

Now consider another hyperplane of B that passes through c . The corresponding hyperplane must be orthogonal to every one of these $(n-1)$ hyperplanes that it meets. Therefore it is parallel to one of them and orthogonal to the rest. So the hyperplanes fall into $n-1$ parallelism classes. Each class can be parameterized by a subset Σ_i of $\mathbb{C}$.

At a formal level, we regard i as varying over the set of parallelism classes of hyperplanes incident to c , rather than $\{1, \dots, n-1\}$. Therefore Γ acts naturally on the set of pairs (c, i) . However, we will continue to write $i = 1, \dots, n-1$ in our concrete descriptions. We have proven most of:

Lemma 7.3. *For $i = 1, \dots, n-1$ there are lattices $\Lambda_i \subset \mathbb{C}$ and Λ_i -invariant discrete subsets Σ_i of $\mathbb{C}$, such that the following holds: $\mathcal{H} \cap \partial T_c$ is the union of the fibers over the arrangement*

$$\bigcup_{i=1}^{n-1} (\mathbb{C}^{i-1} \times \Sigma_i \times \mathbb{C}^{n-1-i}) \subseteq \mathbb{C}^{n-1}.$$

Proof. The lemma is vacuous if $n = 1$, so we can assume $n \geq 2$. We have already proven everything except the existence of the Λ_i and the discreteness of the Σ_i . The

latter follows from the local finiteness of $\mathcal{H}$. Next, since Γ is a lattice in $\mathrm{PU}(n, 1)$, and c is a cusp of Γ , the Γ -stabilizer of T_c is a lattice in the $\mathrm{PU}(n, 1)$ -stabilizer of T_c , i.e. in $H^{2n-1} \rtimes U(n-1)$ (for example because the image of the horosphere must be compact in the quotient). By a variation on Bieberbach's Theorem, this has a finite index normal subgroup which is a lattice in H^{2n-1} [Aus60].

Some pair of elements of this lattice in H^{2n-1} has non-zero commutator, showing that the lattice contains a nonzero element of Z and hence a lattice in Z . The quotient of the lattice in H^{2n-1} by its intersection with Z is a lattice Λ in $H^{2n-1}/Z = \mathbb{C}^{n-1}$. Translation by Λ preserves each parallelism class, hence acts on each set Σ_i . The action on any given Σ_i is by a discrete subgroup of $\mathbb{C}$, which we write Λ_i . Each Λ_i is a lattice in $\mathbb{C}$, or else Λ could not be a lattice in $\mathbb{C}^{n-1}$. $\square$

It is easy to write down the complement of this arrangement, namely

$$\mathbb{C}^{n-1} - \bigcup_{i=1}^{n-1} (\mathbb{C}^{i-1} \times \Sigma_i \times \mathbb{C}^{n-1-i}) = \prod_{i=1}^{n-1} \mathbb{C}_i^\circ,$$

where $\mathbb{C}_i^\circ = \mathbb{C} - \Sigma_i$.

7.3. Pulling back to the universal cover. Next, we pull all these constructions back to $\hat{T}_{\hat{c}}$. First we observe that $\hat{T}_{\hat{c}}^\circ$ is the universal cover of T_c° . This uses the fact that $T_c^\circ \rightarrow B^\circ$ induces an injection on π_1 , as can be seen by noting that closest point projection provides a one-sided inverse. It follows that $\partial\hat{T}_{\hat{c}}^\circ$ is the universal cover of ∂T_c° . Next, the action of $Z \cong \mathbb{R}$ on ∂T_c° lifts to one on $\partial\hat{T}_{\hat{c}}^\circ$. (Pull back the vector field describing the former action, and then integrate it.) This makes $\partial\hat{T}_{\hat{c}}^\circ$ into a principal $\mathbb{R}$ -bundle. Since its total space is simply connected, its base is too (use the long exact sequence of the fiber bundle).

Writing $\hat{\mathbb{C}}_i^\circ$ for the universal cover of $\mathbb{C}_i^\circ$, it follows that $\partial\hat{T}_{\hat{c}}^\circ$ is the pullback of $\partial T_c^\circ \rightarrow \prod \mathbb{C}_i^\circ$ to the universal cover $\prod \hat{\mathbb{C}}_i^\circ$ of its base.

7.4. Tree domains. We can now define the tree domains. For each i , write $\hat{\mathbb{C}}_i$ for the metric completion of $\hat{\mathbb{C}}_i^\circ$. Note that $\hat{\mathbb{C}}_i$ is CAT(0).

Each i , viewed as a parallelism class, is a tree domain; thus we get $n-1$ tree domains from each cusp $\hat{c}$. The associated metric space will be $\mathcal{C}(i) = \hat{\mathbb{C}}_i$.

The projection $\partial\hat{T}_{\hat{c}}^\circ \rightarrow \prod \hat{\mathbb{C}}_i^\circ$ extends to a map $\partial\hat{T}_{\hat{c}} \rightarrow \prod \hat{\mathbb{C}}_i$. The action of $\mathbb{R}$ extends to $\partial\hat{T}_{\hat{c}}$, and the quotient of this action is $\prod \hat{\mathbb{C}}_i$.

The HHS projection $\pi_{\hat{c},i} : \hat{B}^\theta \rightarrow \mathcal{C}(i)$ is the inclusion $\hat{B}^\theta \rightarrow \hat{B}$, followed by closest point projection to $\hat{T}_{\hat{c}}$, followed by the projection map $\partial\hat{T}_{\hat{c}} \rightarrow \prod \hat{\mathbb{C}}_i$, followed by projection to the i th factor.

Remark 7.4. Continuing Remark 4.3, we note that the true domain of this map is $\hat{B}^{\hat{c},i} = \hat{B} - \mathrm{int}(\hat{T}_{\hat{c}})$.

The next lemma shows that $\mathcal{C}(i)$ is Gromov hyperbolic; its projection $\pi_{\hat{c},i}$ is coarsely Lipschitz because it is a composition of coarsely Lipschitz maps.

Lemma 7.5. *Suppose $\hat{c}$ is a cusp of $\hat{B}$, and suppose i is a parallelism class of hyperplanes of $\hat{B}$ that are incident to $\hat{c}$. Then the tree domain $\mathcal{C}(i)$ is Gromov hyperbolic.*

We remark that the parallelism classes here have the obvious meaning: two hyperplanes of $\widehat{B}$, that are incident at $\widehat{c}$, are parallel just if their images in B are parallel as hyperplanes containing c . In particular, parallel hyperplanes are disjoint, and non-parallel hyperplanes meet.

Before we prove the lemma, we set up some of the notation that will also be used later. Let G_i be the graph in $\mathbb{C}_i^\circ$ which is the union of the boundaries of the Voronoi cells around the points of Σ_i . The size of the cells is uniformly bounded, by the Λ_i -invariance of Σ_i .

We write $\widehat{G}_i$ for the preimage of $G_i \subseteq \mathbb{C}_i^\circ$ in $\widehat{\mathbb{C}}_i^\circ$. Because $G_i \rightarrow \mathbb{C}_i^\circ$ is a homotopy equivalence, $\widehat{G}_i$ is the universal cover of G_i . In particular, it is a tree.

Proof. As a tree, $\widehat{G}_i$ is Gromov hyperbolic. By a cell of $\widehat{\mathbb{C}}_i$ we mean a component of the preimage of a Voronoi cell in $\mathbb{C}_i$. Its boundary is a copy of $\mathbb{R}$ in the tree $\widehat{G}_i$. So [Proposition 2.3](#) establishes that the electrification $\widehat{G}_i^!$ of $\widehat{G}_i$ along these copies of $\mathbb{R}$ is Gromov hyperbolic. Thus it suffices to show that $\widehat{\mathbb{C}}_i$ is quasi-isometric to $\widehat{G}_i^!$.

There is a natural map $\widehat{\mathbb{C}}_i \rightarrow \widehat{G}_i^!$ that is the identity on $\widehat{G}_i$ and maps the interior of each cell to the corresponding generic point of $\widehat{G}_i^!$. There is also a natural map $\widehat{G}_i^! \rightarrow \widehat{\mathbb{C}}_i$ that again is the identity on $\widehat{G}_i$ and maps each generic point to the corresponding singular point of $\widehat{\mathbb{C}}_i$. These maps are coarsely inverse to each other and coarsely Lipschitz, proving that they are quasi-isometries. $\square$

The proof of [Lemma 7.5](#) explains the terminology “tree domains”, and gives an alternative way to think of the associated hyperbolic spaces as electrifications of trees.

7.5. Center domains. We view $\partial\widehat{T}_{\widehat{c}}$ as a type of $\mathbb{R}$ -bundle over $\prod \widehat{\mathbb{C}}_i$. The center domain associated to $\widehat{T}_{\widehat{c}}$ will be defined as a swaddling of $\partial\widehat{T}_{\widehat{c}}$. We will use [Lemma 6.4](#) to see that this swaddling is a quasi-line, and it will serve to measure distance in the center direction in the horoball.

The projection map to this quasi-line is defined as the composition of the inclusion $\widehat{B}^\theta \rightarrow \widehat{B}$, followed by nearest-point projection to $\partial\widehat{T}_{\widehat{c}}$, followed by the natural map $\partial\widehat{T}_{\widehat{c}} \rightarrow \partial\widehat{T}_{\widehat{c}}^{\text{sw}}$. This will be coarsely Lipschitz because it is a composition of coarsely Lipschitz maps, given that [Lemma 7.8](#) below shows that $\partial\widehat{T}_{\widehat{c}} \rightarrow \partial\widehat{T}_{\widehat{c}}^{\text{sw}}$ is coarsely Lipschitz.

Remark 7.6. Continuing [Remark 4.3](#), we note that the true domain of this map is $\widehat{B}^{\widehat{T}_{\widehat{c}}} = \widehat{B} - \text{int}(\widehat{T}_{\widehat{c}})$.

Remark 7.7 (Optional motivation). Before we continue, we want to discuss the intuition and some ideas that encounter technical complications, to better motivate how we proceed.

The base $\prod \widehat{\mathbb{C}}_i$ has a sort of spine, defined as $\prod \widehat{G}_i$. This product of trees can be viewed as a type of Lagrangian cube complex. Because it is Lagrangian, one can show that the restriction of $\partial\widehat{T}_{\widehat{c}} \rightarrow \prod \widehat{\mathbb{C}}_i$ becomes, in a certain metric sense, a trivial bundle, so the preimage of $\prod \widehat{G}_i$ can be viewed as $\mathbb{R} \times \prod \widehat{G}_i$. This stands in sharp contrast to the extremely non-trivial bundle $\partial T_c \rightarrow \mathbb{C}^{n-1}$.

Although it is tempting to use a deformation retract $\prod \widehat{\mathbb{C}}_i^\circ$ to $\prod \widehat{G}_i$ to make use of the triviality over the spine, one quickly encounters technical difficulties related to the non-Lipschitz nature of this retract and the fact that bounded diameter loops

in $\prod \hat{\mathbb{C}}_i^\circ$ can enclose arbitrarily large area. For this reason we abandon the spine and work with all of $\partial \hat{T}_\epsilon$.

In the most naive approach to working with all of $\partial \hat{T}_\epsilon$, one might consider paths in the base $\prod \hat{\mathbb{C}}_i$ that are geodesics in each factor. However, despite the fact that each $\hat{\mathbb{C}}_i$ is both CAT(0) and Gromov hyperbolic, geodesic triangles in $\hat{\mathbb{C}}_i$ can enclose arbitrarily large area.

In conclusion, pushing entirely onto the spine causes problems, and not pushing at all onto the spine causes problems, and this motivates the hybrid approach we follow now.

Swaddling $\partial \hat{T}_\epsilon$ requires choosing a family of $\mathbb{R}$ -torsor isomorphisms between pairs of fibers. We do this by defining the *swaddling path* τ_{xy} for any pair of points

$$x = (x_1, \dots, x_{n-1}), y = (y_1, \dots, y_{n-1}) \in \prod \hat{\mathbb{C}}_i.$$

It is the path $(\tau^{(1)}, \dots, \tau^{(n-1)})$, where $\tau^{(i)}$ is the path in $\hat{\mathbb{C}}_i$ defined as follows (see Figure 2). Begin with the geodesic $\overline{x_i y_i}$. If it enters and leaves a cell C of $\hat{\mathbb{C}}_i$, at points of ∂C whose angular separation is less than 2π (from the perspective of the center of C), then we replace that portion of $\overline{x_i y_i}$ by the portion of ∂C between them.




FIGURE 2. The definition of $\tau^{(i)}$.

Doing this simultaneously, for every cell C , yields a piecewise geodesic path in $\hat{\mathbb{C}}_i$. We define $\tau^{(i)}$ as the parameterization of this path by $[0, 1]$, proportionally to arclength.

One can almost think of it as a path in the tree $\hat{G}_i$, except that it will enter the cells containing its endpoints, and will cut corners, through the centers of cells, whenever this provides a “significant” shortcut.

Let q denote the natural map $\prod \hat{\mathbb{C}}_i \rightarrow \prod \mathbb{C}_i$. Each swaddling path projects to a concatenation of line segments in $\prod \mathbb{C}_i$. Consider each of these segments σ in turn. The fibers of $\partial \hat{T}_\epsilon$ over its endpoints are identified with the fibers of ∂T_c over the endpoints of $q(\sigma) \subseteq \prod \mathbb{C}_i$, since $\partial \hat{T}_\epsilon \rightarrow \prod \hat{\mathbb{C}}_i$ is the pullback of $\partial T_c \rightarrow \prod \mathbb{C}_i$. Furthermore, the Ehresmann connection on $\partial T_c \rightarrow \prod \mathbb{C}_i$ trivializes over $q(\sigma)$. This provides an identification of the fibers of ∂T_c over the endpoints of $q(\sigma)$, and therefore an identification of the fibers of $\partial \hat{T}_\epsilon$ over the endpoints of σ . Concatenating these identifications, as σ varies over the segments comprising the swaddling path, yields an $\mathbb{R}$ -torsor isomorphism from the fiber of $\partial \hat{T}_\epsilon$ over x to the fiber over y .

We apply the swaddling construction to the family of $\mathbb{R}$ -torsor isomorphisms induced by all swaddling paths, obtaining the swaddled space $\partial \hat{T}_\epsilon^{\text{sw}}$. The horoball gives a center domain, which we also call $\hat{T}_\epsilon$, and we define its associated metric space $\mathcal{C}(\hat{T}_\epsilon)$ to be $\partial \hat{T}_\epsilon^{\text{sw}}$.

Lemma 7.8. *The inclusion $\partial\hat{T}_{\hat{c}} \rightarrow \partial\hat{T}_{\hat{c}}^{\text{sw}}$ is coarsely Lipschitz.*

Proof. We provide a sketch only.

Consider two points of $\partial\hat{T}_{\hat{c}}^{\circ}$ joined by a path α in $\partial\hat{T}_{\hat{c}}^{\circ}$ of length at most 1. Assume α is parametrized by $[0, 1]$. By an approximation argument and Lemma 6.5, it suffices to show that $\alpha(1)$ is a uniformly bounded central translation of $\hat{m}_{x,y}(\alpha(0))$, where x and y are the images of $\alpha(0)$ and $\alpha(1)$ in $\prod_i \hat{\mathbb{C}}_i$.

The connection 1-form of $T_c \rightarrow \mathbb{C}^{n-1}$ is bounded because it is invariant under the stabilizer of T_c . Hence the connection 1-form of $\hat{T}_{\hat{c}}^{\circ} \rightarrow \prod \hat{\mathbb{C}}_i^{\circ}$ is also bounded. Let β be the horizontal lift, starting at $\alpha(0)$, of the image of α in $\prod \hat{\mathbb{C}}_i$. Note that $\beta(1)$ and $\alpha(1)$ are in the same fiber of the map to $\prod \hat{\mathbb{C}}_i$, and hence differ by translation of some $t \in \mathbb{R}$. We note that t is bounded by a constant times the length of α , because it is the integral over α of the connection 1-form.

It now suffices to show that $\beta(1)$ is a uniformly bounded central translation of $\hat{m}_{x,y}(\alpha(0))$. Consider the image of β in a factor $\hat{\mathbb{C}}_i$, and consider also the geodesic in $\hat{\mathbb{C}}_i$ joining its endpoints. This loop encloses bounded area in $\hat{\mathbb{C}}_i$, for example by [BH99, page 426]. Additionally, the $\hat{\mathbb{C}}_i$ geodesic and the $\tau^{(i)}$ path enclose bounded area in $\hat{\mathbb{C}}_i$, because of the definition of the τ path and the fact that a bounded length geodesic in a $\hat{\mathbb{C}}_i$ can only go through boundedly many Voronoi cells. The result follows because horizontal lifts along different paths differ by a central translation proportional to the sum of the signed areas enclosed in the $\hat{\mathbb{C}}_i$. $\square$

Theorem 7.9. *For any cusp $\hat{c}$ of $\hat{B}$, the metric space $\mathcal{C}(\hat{T}_{\hat{c}}) = \partial\hat{T}_{\hat{c}}^{\text{sw}}$ is a quasi-line and hence Gromov hyperbolic.*

We write A for the largest of the areas of the Voronoi cells of $\mathbb{C}_i$, over all i .

Lemma 7.10. *Suppose $x, z \in \hat{\mathbb{C}}_i$ and $b \in \overline{xz}$. Then the area enclosed by the triangle $\tau_{xb}^{(i)} \tau_{bz}^{(i)} \tau_{zx}^{(i)}$ is at most A .*

Proof. The essential case is when there is a cell C of $\hat{\mathbb{C}}_i$, whose interior contains b but not x or z . See Figure 3.




FIGURE 3. The proof of Lemma 7.10

We define a and c as the points where ∂C meets $\overline{xb}$ and $\overline{bz}$ respectively. Then $\tau_{xz}^{(i)}$ coincides with $\tau_{xb}^{(i)}$ between x and a , and with $\tau_{bz}^{(i)}$ between c and z . So $\tau_{xb}^{(i)} \tau_{bz}^{(i)} \tau_{zx}^{(i)}$ encloses the same area as $\tau_{ab}^{(i)} \tau_{bc}^{(i)} \tau_{ca}^{(i)}$.

Now, $\tau_{ab}^{(i)}$ is $\overline{ab}$, $\tau_{bc}^{(i)}$ is $\overline{bc}$, and $\tau_{ac}^{(i)}$ is either $\overline{ac}$ (if the angular separation between a and c is $\geq 2\pi$) or the portion of ∂C between a and c (otherwise). In the former case, $\tau_{ac}^{(i)}$ coincides with $\tau_{ab}^{(i)} \tau_{bc}^{(i)}$, so the triangle they form encloses no area. In the latter case, the triangle they form is bounded by ∂C and $\overline{ab} \cup \overline{bc} = \overline{ac}$. Because the

angular separation between a and c is $< 2\pi$, the area of this region is at most that of the Voronoi cell of $\mathbb{C}_i$ over which C lies. So it is at most A .

The argument simplifies in the various degenerate cases. Namely, if x resp. z lies in the interior of C , then set $a = x$ resp. $c = z$, and modify the argument slightly. The other possible degeneration is when b does not lie in the interior of any cell. Then it lies in $\hat{G}_i$, and $\tau_{xz}^{(i)} = \tau_{xb}^{(i)} \tau_{bz}^{(i)}$, so the triangle encloses no area. $\square$

Lemma 7.11. *Suppose $x, y, z \in \hat{\mathbb{C}}_i$, and $a \in \overline{xz} - \{x, z\}$, such that $\overline{xa}$ and $\overline{za}$ make angle $\geq 2\pi$ at a . Then either $\overline{xa} \subseteq \overline{xy}$ or $\overline{za} \subseteq \overline{zy}$.*

Proof. Because $\hat{\mathbb{C}}_i^\circ$ is locally Euclidean, two geodesics whose union is a geodesic make angle π . Therefore $a \notin \hat{\mathbb{C}}_i^\circ$, so a is a ramification point of $\hat{\mathbb{C}}_i$. The rays emanating from a form a copy of $\mathbb{R}$, and $\overline{ax}$ and $\overline{az}$ represent two of these rays.

Let R be their angle bisector, i.e. the geodesic emanating from a , whose angles with $\overline{ax}$ and $\overline{az}$ are equal. It might happen that R eventually hits another branch point of $\hat{\mathbb{C}}_i$ (in which case we stop it there, so R is a segment), or not (in which case R is a ray).

In either case, R cuts $\hat{\mathbb{C}}_i$ into two components, one containing x and the other z . Therefore R meets one of the other edges $\overline{xy}, \overline{yz}$ of the triangle xyz . Without loss of generality, suppose it meets $\overline{yz}$. See Figure 4. Then the angle at a , between





FIGURE 4. The proof of Lemma 7.11

$\overline{ay}$ and $\overline{az}$, is at least as large as the angle between R and $\overline{az}$, which is at least π . Therefore their concatenation is a geodesic, in fact the geodesic $\overline{yz}$, which therefore contains $\overline{az}$. $\square$

Proof of Theorem 7.9. We must show that the monodromy, around any triangle of swaddling paths in $\prod \hat{\mathbb{C}}_i$, is uniformly bounded. Because of the product structure, it is enough to show this for each $\hat{\mathbb{C}}_i$ separately. We claim that for any $x, y, z \in \hat{\mathbb{C}}_i$, the area enclosed by $\tau_{xy}^{(i)} \tau_{yz}^{(i)} \tau_{zx}^{(i)}$ is at most $9A$. We define x' as the point of $\overline{xy} \cap \overline{xz}$ furthest from x , and similarly for y' and z' . See Figure 5. By construction, any pair of edges, of the geodesic triangle $x'y'z'$, meet only at their common vertex. Also, the points x, x', y', y occur in this order along $\overline{xy}$ (possibly with coincidences), and similarly for $\overline{yz}$ and $\overline{zx}$.

By six applications of Lemma 7.10, the signed area enclosed by $\tau_{xy}^{(i)} \tau_{yz}^{(i)} \tau_{zx}^{(i)}$ differs by $\leq 6A$ from the signed area enclosed by the nonagon

$$\tau_{xx'}^{(i)} \tau_{x'y'}^{(i)} \tau_{y'y}^{(i)} \cdot \tau_{yy'}^{(i)} \tau_{y'z'}^{(i)} \tau_{z'z}^{(i)} \cdot \tau_{zz'}^{(i)} \tau_{z'x'}^{(i)} \tau_{x'x}^{(i)}.$$


FIGURE 5. The definition of x', y', z'

This nonagon is the triangle $\tau_{x'y'}^{(i)}\tau_{y'z'}^{(i)}\tau_{z'x'}^{(i)}$ with bigons $\tau_{x'x}^{(i)}\tau_{xx'}^{(i)}$, $\tau_{y'y}^{(i)}\tau_{yy'}^{(i)}$ and $\tau_{z'z}^{(i)}\tau_{zz'}^{(i)}$ attached at its corners. Each bigon encloses zero area, because the swaddling paths comprising it are each other's reversals. So it is enough to show that $\tau_{x'y'}^{(i)}\tau_{y'z'}^{(i)}\tau_{z'x'}^{(i)}$ encloses area $\leq 3A$.

Having no further need for x, y, z , we will suppress the primes from x', y', z' . Repeating a known fact in this notation: $(\star)$ any pair of edges, of the geodesic triangle xyz , meet only at their common vertex. If two of x, y, z coincide, then $(\star)$ forces all three to coincide. In this case, $\tau_{xy}^{(i)}\tau_{yz}^{(i)}\tau_{zx}^{(i)}$ degenerates to a point, so it encloses no area. So we may suppose the three vertices are distinct.

We claim that for every $a \in \overline{xz} - \{x, z\}$, $\overline{ax}$ and $\overline{az}$ make angle $< 2\pi$. Otherwise, [Lemma 7.11](#) would lead to violation of $(\star)$. From this follows a key fact: if $\overline{xz}$ enters and exits a cell C at points of ∂C , then their angle (viewed from the center of C) is $< 2\pi$. Therefore, in the definition of $\tau_{xz}^{(i)}$, the portion of $\overline{xz}$ that passes through C is replaced by a portion of ∂C . Applying the same reasoning to the other cells meeting $\overline{xz}$, and similarly for the other two edges, we conclude: $\tau_{xy}^{(i)}\tau_{yz}^{(i)}\tau_{zx}^{(i)}$ encloses no portion of any cell, except perhaps for cells whose interiors contain at least one of x, y or z . So it suffices to show that the enclosed portion of such a cell has area $\leq A$.

We examine such a cell C , writing c for its center. If one of the vertices is c , say $x = c$, then consider the geodesics from x to ∂C , in the directions of $\overline{xy}$ and $\overline{xz}$. For this and the next cases, see [Figure 6](#), starting on the left for this case. These make

FIGURE 6. Three cases for the conclusion of the proof of [Theorem 7.9](#). In all cases the shaded region has area at most A .

angle less than π ; otherwise, $\overline{xy} \cup \overline{xz}$ would be a geodesic, namely $\overline{yz}$, which would contradict $(\star)$. The portion of C enclosed by $\tau_{xy}^{(i)} \tau_{yz}^{(i)} \tau_{zx}^{(i)}$ lies in the region bounded by these two geodesics and ∂C . Since their angle is less than π , this region's area is $\leq A$.

Now suppose c is distinct from x, y, z , but lies in one of the edges of the geodesic triangle, say $\overline{xz}$. The angle between $\overline{cx}$ and $\overline{cz}$ is less than 2π , or otherwise [Lemma 7.11](#) would yield a contradiction of $(\star)$. Consider the geodesics from c to ∂C in these two directions, and argue as in the previous paragraph.

Finally suppose no edge of the triangle contains c . Of the edges of the triangle, choose one which comes closest to c , say $\overline{xy}$. If it is “radial”, meaning that it extends to a geodesic that hits c , then its end near c lies in another edge of the triangle, which is non-radial and comes as close to c as $\overline{xy}$ does. Using it instead, we may suppose without loss that $\overline{xy}$ is not radial. Being non-radial, $\overline{xy} \cap C$ extends to a geodesic between two points of ∂C , that misses c . The portion of C enclosed by $\tau_{xy}^{(i)} \tau_{yz}^{(i)} \tau_{zx}^{(i)}$ lies in the region bounded by this geodesic and ∂C , which has area $\leq A$. $\square$

7.6. Enough projections. We end this section with a statement that will be used later in the proof of the enough projections assumption of [Proposition 3.6](#).

Lemma 7.12. *For each $\kappa \geq 0$ there exists $\omega \geq 0$ such that if $x, y \in \partial \hat{T}_{\hat{c}}$ and $d(x, y) \geq \omega$ then $d_U(\pi_U(x), \pi_U(y)) \geq \kappa$ for some cusp domain U associated to $\hat{T}_{\hat{c}}$.*

Here we use that the π_U are in fact defined on all of $\partial \hat{T}_{\hat{c}}$, because $\hat{B}^U$ is the complement of the interior of $\hat{T}_{\hat{c}}$. Note that it ultimately does not matter if $d(x, y)$ denotes the $\hat{B}$ metric or the intrinsic path metric on $\partial \hat{T}_{\hat{c}}$ here because of [Lemma 8.18](#), but for concreteness let us say it denotes the later for now.

Proof. We make use of the map $\partial \hat{T}_{\hat{c}} \rightarrow \prod \hat{C}_i$. Assume $x, y \in \partial \hat{T}_{\hat{c}}$ have that $d_U(\pi_U(x), \pi_U(y)) \leq \kappa$ for all cusp domains U associated to $\hat{T}_{\hat{c}}$. We must give some upper bound ω on $d(x, y)$.

Since the images of x and y in each factor $\hat{C}_i$ are close, the length of the swaddling path from x to y is bounded. We can use the swaddling data above to “parallel translate” x to a point x' in the same fiber of y . The swaddling data ultimately comes from horizontal lifts, so the distance from x to x' is bounded. Thus it suffices to bound the distance from x' to y .

But x' and y are in the same fiber, and their distance in that fiber is bounded by [Lemma 6.4](#) and the assumption. $\square$

8. ANGLE DOMAINS

The angle domains will be in bijection with the hyperplanes $\hat{H}$ in $\hat{B}$, but to avoid confusion with the associated strata we define them to be indexed by the unit normal bundles $N^1(\hat{H})$ to the hyperplanes.

The idea of our angle domains is simple, but takes work to set up. There is a function on $\mathbb{R}^3 - \mathbb{R}$ which indicates angular position around $\mathbb{R}$. Usually one thinks of it as taking values in a circle. This function lifts to one on the universal cover of $\mathbb{R}^3 - \mathbb{R}$, taking values in a copy of $\mathbb{R}$. Our goal is to construct a similar function on $\hat{B} - \hat{H}$, for each hyperplane $\hat{H} \subseteq \hat{B}$. Our “angle” function will not take values in a copy of $\mathbb{R}$, but rather in a quasi-line, i.e. a space quasi-isometric to $\mathbb{R}$. The

hyperbolic spaces associated to angle domains in our HHS structure will be the quasi-lines associated to hyperplanes of $\hat{B}$, and the projections to them will be their “angle” functions. The quasi-lines are defined intrinsically. In particular, the $\hat{\Gamma}$ -stabilizer of $\hat{H}$ acts on $\hat{H}$ ’s quasi-line.

Remark 8.1. These domains are analogous to annular domains for mapping class groups, and they can also be compared, roughly speaking, to a Fenchel-Nielsen twist parameter for Teichmüller space. In the context of Teichmüller space one observes that attempts to measure twisting or annular subsurface projections are subtle and sometimes difficult, and there are non-equivalent definitions with different advantages (see [AHW25, Appendix A] for some discussion). Experts on hierarchical hyperbolicity will also recall [Man25] as a compelling explanation that projections to annular curve graphs are not at all as canonical as one might have initially expected. Thus, even though the hyperbolic spaces for our angle domains seem very canonical, it is not a surprise that there is more than one way to define the projections to them; and in fact we exploit this flexibility to choose a definition which avoids some technical issues.

8.1. Unit normal bundles. Fix a hyperplane $\hat{H}$ of $\hat{B}$ and write $p_{\hat{H}}$ for nearest-point projection to it, which makes sense because $\hat{B}$ is $\text{CAT}(-1)$ and $\hat{H}$ is convex. Suppose $\ell > 0$ and $\hat{\gamma} : [0, \ell] \rightarrow \hat{B}$ is a geodesic with $\hat{\gamma}(0) \in \hat{H}$. We say $\hat{\gamma}$ is orthogonal to $\hat{H}$ if $p_{\hat{H}} \circ \hat{\gamma}(t)$ is constant. We define the unit normal bundle $N^1(\hat{H})$ to be the set of germs of such geodesics, meaning that two geodesics are considered equivalent if they coincide on some interval $[0, \varepsilon]$ with $\varepsilon > 0$. There is a surjective projection map

$$p_{N^1(\hat{H})} : \hat{B} - \hat{H} \rightarrow N^1(\hat{H}),$$

sending each $\hat{y}$ to the germ of the geodesic $\overline{p_{\hat{H}}(\hat{y})\hat{y}}$.

The pointwise stabilizer of $\hat{H}$ in $\pi_1(B^\circ) \subset \hat{\Gamma}$ is a subgroup isomorphic to $\mathbb{Z}$. By covering space theory (applied after removing hyperplanes and before taking metric completions), this corresponds to a branched cover $B' \rightarrow B$ with deck group $\mathbb{Z}$. Let H' be the image of $\hat{H}$ in B' , and as usual let H be the image of $\hat{H}$ in B . Note that $B' - H'$ is the universal cover of $B - H$.

Exactly the same considerations as above apply to the hyperplane H' of B' , yielding the unit normal bundle $N^1(H')$ of H' in B' . The pointwise stabilizer of H in $\text{Isom}^+(B)$ is a copy of S^1 , which rotates around H . This S^1 action on B lifts via covering space theory to a $\mathbb{R}$ action on B' which fixes H' pointwise. The induced action on $N^1(H')$ shows that $N^1(H')$ is a principal $\mathbb{R}$ -bundle. Our goal in this subsection is to show the following.

Lemma 8.2. *$N^1(\hat{H})$ is a $\mathbb{R}$ -bundle, and is the pullback of $N^1(H')$ via the map $\hat{H} \rightarrow H'$.*

Proof. Let H be the image of $\hat{H}$ in B .

For each $p \in H$ there is some ϵ such that any hyperplane coming within ϵ of p in fact contains p . If y is within ϵ of p , then the geodesic from y to H lies in a single stratum (except the endpoint in H). Any such geodesic segment can be lifted to $\hat{B}$ and B' to geodesics that end on $\hat{H}$ or H' .

All germs in $N^1(\hat{H})$ and $N^1(H')$ arise in this way. This analysis shows in particular that there is a surjective map from $N^1(\hat{H})$ to the pullback of $N^1(H')$. This

map sends the germ of a geodesic at a point of $\hat{H}$ to the pair of that point of $\hat{H}$ and the pushforward germ of a geodesic based at the image of that point in H' .

We can see this map is injective as follows. If not, it must be because there are two germs w, w' in $N^1(\hat{H})$ with the same basepoint $x \in \hat{H}$ that map to the same point of $N^1(H')$. In particular, they must differ by an element γ of the deck group of $\hat{B} \rightarrow B'$ that fixes x . If x lies in k hyperplanes, its stabilizer in the subgroup of $\pi_1(B^\circ) \subset \hat{\Gamma}$ corresponding to B is a $\mathbb{Z}^k$. The different coordinates of $\mathbb{Z}^k$ correspond to different hyperplanes containing x . However, as can be seen using B , every hyperplane through x other than $\hat{H}$ contains all germs of geodesics based at x and orthogonal to $\hat{H}$. So actually w must be fixed by the $\mathbb{Z}^{k-1}$ corresponding to the $k-1$ hyperplanes other than $\hat{H}$. Applying the $\mathbb{Z}$ corresponding to $\hat{H}$ produces germs with different image in $N^1(H')$, so this gives the result.

The topology and principal $\mathbb{R}$ -bundle structure on $N^1(\hat{H})$ can now be obtained by pulling back. $\square$

The map from H' to H is in fact an isomorphism, and we can think of H' as a copy of $\mathbb{CH}^{n-1}$.

Equip $N^1(H)$ with the Levi-Civita connection. Viewed as a 1-form on $N^1(H)$ with values in $\mathbb{R}$ (the Lie algebra of S^1), it can be pulled back to a connection on $N^1(H')$, via the natural map $N^1(H') \rightarrow N^1(H)$.

The isometry group of B' fixes H' as a set and acts transitively on H' . The curvature form of $N^1(H')$ is invariant under this action.

Remark 8.3. In fact, symmetry shows that the curvature 2-form is a multiple of the imaginary part of the Hermitian form on the hyperplane, and it is also possible to compute the curvature form directly using [KN63, pages 77, 301]. There is also a soft argument to show that the curvature cannot be zero, as a consequence of a certain central extension $P(U(1, n-1) \times S^1) \rightarrow PU(1, n-1)$ not splitting.

8.2. Preparatory lemmas. We also note the following at the level of B , where we also have a map $p_{N^1(H)} : B - H \rightarrow N^1(H)$.

Lemma 8.4. *Suppose $\alpha(t), t \in [0, 1]$ is a path with $\alpha(0) \in H$ and $\alpha'(0)$ non-zero and orthogonal to H . Then*

$$\lim_{t \rightarrow 0^+} p_{N^1(H)}(\alpha(t)) = \frac{\alpha'(0)}{\|\alpha'(0)\|}.$$

For a proof, compare to [Lee18, Theorem 5.25, Proposition 5.26, Corollary 6.39].

8.3. The hyperbolic spaces. We start by recalling the following.

Lemma 8.5. *Let η be a closed 2-form on $\mathbb{CH}^{n-1}$ which is invariant under the group of holomorphic isometries of $\mathbb{CH}^{n-1}$. Then there exists L such that for any geodesic triangle T in $\mathbb{CH}^{n-1}$, the absolute value of the integral of η over any filling of T is at most L times the area of the comparison triangle.*

Proof. Because η is invariant under a transitive action by isometries, it is bounded. Recall that any geodesic triangle in complex hyperbolic space can be filled by a disc whose area is at most the area of the comparison triangle in $\mathbb{H}^2$ [BH99, page 426] (and since the exponential map is smooth, the disc can be taken to be smooth on its interior). The result follows. $\square$

Parallel translation along geodesics defines geometric swaddling data for $N^1(H')$. Since holonomy around a triangle can be computed as the integral of curvature on a filling disc (see for example [MMR90, page 41]), keeping in mind [Definition 6.9](#) we get the following.

Corollary 8.6. *The geometric swaddling data for $N^1(H')$ is L -bounded for some L .*

We now wish to pull back from H' to $\hat{H}$. To that end, recall the definition of a piecewise triangular map from [Definition 6.18](#), and note the following.

Lemma 8.7. *Let S be a ball, and $\hat{S}$ the metric completion of the universal cover of an orthogonal arrangement in S . Then the map $\hat{S} \rightarrow S$ is piecewise triangular.*

It follows in particular that it is piecewise geodesic in the sense of [Definition 6.15](#).

Proof. To start with, consider an open cover of $\hat{S}$ consisting of, for each point p of $\hat{S}$, a ball centered at p small enough to be disjoint from all hyperplanes not containing p . [Lemma 6.13](#) gives that it suffices to consider a triangle T in such a ball U . Let Δ denote the intersection of all the hyperplanes through p . We subdivide T as follows, as illustrated in [Figure 7](#) for the first three steps.

- (1) If a vertex a of T lies in Δ , then we subdivide the opposite edge into segments, each of which lies in a single stratum (except perhaps for its endpoints). We subdivide T along the geodesics from a to these points. Each of the resulting triangles has edges that each map to a geodesic in S .




FIGURE 7. The proof of [Lemma 8.7](#)

- (2) If no vertex lies in Δ , but some point w of an edge of T does, then we subdivide T along the geodesic from w to the opposite vertex. The previous case applies to both parts of the subdivision.
- (3) If $\overline{aw}$ meets Δ , for some vertex a of T and w in the opposite edge, then we subdivide T along $\overline{aw}$. The previous case applies to both parts of the subdivision.
- (4) If none of the previous cases apply, then the set W swept out by the $\overline{aw}$ misses Δ , where a is a vertex of T and w varies over its opposite edge. We apply the Alexandrov patchwork argument from the proof of [Lemma 6.13](#). The required neighborhoods, of the points of W , exist by an inductive argument.

This concludes the proof. $\square$

Keeping in mind [Lemma 8.2](#), we get that [Lemma 6.16](#) provides geometric swaddling data for $N^1(\hat{H})$. [Corollary 8.6](#) and [Corollary 6.19](#) thus give that the swaddled space is a quasi-line:

Corollary 8.8. $N^1(\hat{H})^{\text{sw}}$ is a quasi-line.

For the angle domain $U = N^1(\hat{H})$, define $\mathcal{C}(U) = N^1(\hat{H})^{\text{sw}}$. We also note the following.

Lemma 8.9. The inclusion of $N^1(\hat{H})$ into $N^1(\hat{H})^{\text{sw}}$ is coarsely Lipschitz.

Here we equip $N^1(\hat{H})$ with the path metric obtained by pulling back the natural invariant metric on $N^1(H')$ (the lifted Sasaki metric).

Proof. This is left to the reader; it is very similar to, but easier than, [Lemma 7.8](#). $\square$

Lemma 8.10. Let $\hat{H}_1, \dots, \hat{H}_k$ be hyperplanes that all intersect, and let $U_1, \dots, U_k$ be the associated angle domains. Then any orbit map for the action of $\mathbb{Z}^k$ on $\mathcal{C}(U_1) \times \dots \times \mathcal{C}(U_k)$ gives a quasi-isometry from $\mathbb{Z}^k$ to $\mathcal{C}(U_1) \times \dots \times \mathcal{C}(U_k)$, with uniform constants.

Let x be a generic point of the intersection of the $\hat{H}_i$, i.e. $x \in (\cap_{i=1}^k \hat{H}_i)^\circ$. The $\mathbb{Z}^k$ is the stabilizer of any such x for the deck group of $\hat{B}^\circ \rightarrow B^\circ$ and is described in [Section 5.1](#).

Proof. Pick x as above. For each $1 \leq i \leq k$, the fiber of $N^1(\hat{H}_i)$ at x is, by definition, a copy of $\mathbb{R}$. The result follows from the fact that the $\mathbb{Z}^k$ quotient of the product of these copies of $\mathbb{R}$ is a compact torus (see the discussion in [Section 5.1](#)). $\square$

8.4. The naive projection map. We start our discussion with the map

$$p_{N^1(\hat{H})} : \hat{B} - \hat{H} \rightarrow N^1(\hat{H}).$$

Given a point x of $\hat{B} - \hat{H}$, we consider its closest point projection $p_{\hat{H}}(x)$ to $\hat{H}$, and we recall that by definition $p_{N^1(\hat{H})}(x)$ is the germ of the geodesic from $p_{\hat{H}}(x)$ to x .

This map is not coarsely Lipschitz, nor is its composition with the inclusion $N^1(\hat{H}) \rightarrow N^1(\hat{H})^{\text{sw}}$. In fact, any tiny neighborhood of any point of $\hat{H}$, minus $\hat{H}$, will map to a coarsely dense subset of $N^1(\hat{H})^{\text{sw}}$. It is possible, with perhaps more effort than one would expect, to show that when one restricts to $\hat{B}^\theta$ the map to $N^1(\hat{H})^{\text{sw}}$ is coarsely Lipschitz, but for more than one reason we will choose to circumvent the required analysis for this. What is easy to prove is the following.

Lemma 8.11. For every $C, \epsilon > 0$ there exists a D such that the following is true. Suppose $x, x' \in \hat{B}$ are joined by a path in $\hat{B}$ of length at most 1 that does not come within ϵ of $\hat{H}$. Suppose that $p_{\hat{H}}(x)$ and $p_{\hat{H}}(x')$ have distance at most C from the complement of all horoballs in $\hat{B}$. Then $p_{N^1(\hat{H})}(x)$ and $p_{N^1(\hat{H})}(x')$ have distance at most D in $N^1(\hat{H})^{\text{sw}}$.

The same result applies if x, x' are joined by a geodesic of any length that does not come within ϵ of $\hat{H}$.

Here we assume that the projections to $\hat{H}$ are not deep in a horoball, and we are implicitly using that $N^1(\hat{H})$ is naturally a subset of $N^1(\hat{H})^{\text{sw}}$, and we are following our convention that whenever we refer to a horoball of $\hat{B}$ we mean a lift of one of the carefully chosen horoballs T_c of B , which are based at cusps of Γ and are chosen in particular to be disjoint from each other.

Proof. Pick $0 < a < \min(\epsilon, r_\theta)$ such that in the set of points at most distance $C + 2$ from the complement of the horoballs, disjoint hyperplanes have distance at least $2a$ from each other. Consider the projection of x and x' to the a neighborhood of $\hat{H}$. Since this neighborhood is convex, the projection of the path has length at most 1.

Along this path, the $p_{N^1(\hat{H})}$ image can only change at uniformly bounded speed, because the path stays exactly a away from $\hat{H}$. Indeed, $p_{N^1(H)}$ is uniformly Lipschitz on the set of points exactly distance a from H , and one can lift this using Lemma 8.2 to obtain the corresponding fact for $p_{N^1(\hat{H})}$ as long as one restricts to points only boundedly deep in horoballs.

It follows that $p_{N^1(\hat{H})}(x)$ and $p_{N^1(\hat{H})}(x')$ have uniformly bounded distance from each other in $N^1(\hat{H})$, and Lemma 8.9 gives the first result.

The second result can be derived from the first using Lemma 2.7. $\square$

8.5. A different projection defined on the boundary of a horoball. The rest of this section requires the setup from Section 7. The naive projection above rests on the most natural path from a point $x \in \hat{B} - \hat{H}$ to $\hat{H}$, namely the geodesic to the closest point projection. There is sometimes a different option with subtle technical advantages, and we explore this now.

Let $\hat{T}_\epsilon$ be a horoball that $\hat{H}$ enters, and let x be a point in $\partial\hat{T}_\epsilon - \hat{H}$. There is some i so that $\hat{H}$ corresponds to a point $q_{\hat{H}}$ in $\hat{\mathbb{C}}_i - \hat{\mathbb{C}}_i^\circ$. Inside of $\prod_j \hat{\mathbb{C}}_j$ we can define a path which starts at the image of x , is constant in all but the i -th coordinate, and is the geodesic to $q_{\hat{H}}$ in the i -th coordinate. We then consider the horizontal lift of this path starting at x . It is a path in $\partial\hat{T}_\epsilon$ from x to a point of $\hat{H} \cap \partial\hat{T}_\epsilon$.

Lemma 7.2 gives that the germ of such a path at its endpoint defines a point of $N^1(\hat{H})$. This is our alternative projection.

Lemma 8.12. *For all $\epsilon > 0$ there exists $C_\epsilon > 0$ such that for any point of $\partial\hat{T}_\epsilon$ whose distance to $\hat{H}$ is in $(0, \epsilon)$, the naive and alternative projections to $N^1(\hat{H})^{\text{sw}}$ are within C_ϵ of each other.*

Proof. We first prove the result for some $\epsilon > 0$, and then deduce from that that it holds for all $\epsilon > 0$.

The image in $\hat{\mathbb{C}}_i$ of the path chosen above consists of segments between singularities, and the final segment is the part of the path after it hits the penultimate singularity and before it gets to $q_{\hat{H}}$. If ϵ is small enough, the whole path is equal to its final segment.

The result for a sufficiently small ϵ follows by a compactness argument, since one can project these paths to B , where up to the group action we get a compact set of paths, and then use Lemma 8.4, keeping in mind Lemma 8.9.

Now, suppose the result is known for a small $\epsilon_0 > 0$, and let $\epsilon > \epsilon_0$ be arbitrary. Lemma 8.11 gives that the naive projection changes by only a bounded amount as one moves on $\partial\hat{T}_\epsilon$ from distance ϵ to distance ϵ_0 from $\hat{H}$. The alternative projection does not change at all if one follows a horizontal lift of a geodesic in $\prod \hat{\mathbb{C}}_i$. This gives the result. $\square$

Remark 8.13. It does not matter if one uses the τ path in the i -th coordinate instead of the geodesic, since they give the same germ at $\hat{H}$.

The alternative projection has the following property.

Proposition 8.14. *There is a constant F such that the following holds. Let $\hat{T}_{\hat{c}}$ be a horoball, let $\hat{\mathbb{C}}_i$ be one of its factors, and let $\hat{H}$ be a hyperplane corresponding to a point $q_{\hat{H}}$ in $\hat{\mathbb{C}}_i$. Let U be the angle domain corresponding to $\hat{H}$. If $x, y \in \partial\hat{T}_{\hat{c}} - \hat{H}$ and the geodesic between their images x_i, y_i in $\hat{\mathbb{C}}_i$ does not pass through $q_{\hat{H}}$, then the alternative projections of x and y to $N^1(\hat{H})^{\text{sw}}$ are within F of each other.*

The proof requires some preparation. Start by noting that a neighborhood of $q_{\hat{H}}$ is isometric to a neighborhood of the special point in the metric completion of the universal cover of $\mathbb{R}^2 - \{(0, 0)\}$. The space of geodesics leaving $q_{\hat{H}}$ is parametrized by $\mathbb{R}$, and there is a natural notion of angle between two geodesics leaving $q_{\hat{H}}$ valued in $[0, \infty)$.

Lemma 8.15. *Suppose w and w' are elements of $N^1(\hat{H})$ lying over possibly different points of $\partial\hat{T}_{\hat{c}}$ that both define the same outgoing germ at $q_{\hat{H}} \in \hat{\mathbb{C}}_i$. Then the distance between the images of w and w' in $N^1(\hat{H})^{\text{sw}}$ is universally bounded.*

Proof. By an approximation argument, it suffices to assume w and w' lie over points of the generic stratum of $\hat{H}$. Join the basepoints of w and w' in $\hat{H} \cap \partial\hat{T}_{\hat{c}}$ with a path α in the generic stratum of $\hat{H} \cap \partial\hat{T}_{\hat{c}}$.

For $s \geq 0$, let $\alpha^{(s)}$ denote the image of α under flow for time s towards $\hat{c}$. Thus, the length of $\alpha^{(s)}$ goes to zero as s goes to infinity. Let w_s and w'_s be the parallel translates of w, w' for time s along the geodesics towards $\hat{c}$.

By definition, in $N^1(\hat{H})^{\text{sw}}$, we know that w_s is distance at most 1 from w and that w'_s is distance at most 1 from w' . So it suffices to show that w_s and w'_s are bounded distance apart in $N^1(\hat{H})^{\text{sw}}$.

A slight variant of [Lemma 7.1](#) gives that w'_s and the parallel translate of w_s along $\alpha^{(s)}$ are very close in $N^1(\hat{H})$ when s is large.

One can consider the loop formed by $\alpha^{(s)}$ and the geodesic between its endpoints. One can approximate $\alpha^{(s)}$ with a piecewise geodesic path with a fine mesh and close to the same length, and consider the geodesic polygon formed by this piecewise geodesic path and the geodesic between the endpoints of $\alpha^{(s)}$. Since this polygon has small perimeter, it can be subdivided into triangles whose total comparison area is small; compare to [\[BH99, page 426\]](#). Because the swaddling data defined on $N^1(\hat{H})$ is L -bounded for some L (by [Corollary 6.19](#)), the total monodromy of the geodesic polygon is small, so by approximation we see that the parallel translate along $\alpha^{(s)}$ is close to the parallel translate along the geodesic between its endpoints. This gives the result. $\square$

Proof of Proposition 8.14. Consider the geodesics from $q_{\hat{H}}$ to x_i and y_i . If the angle at $q_{\hat{H}}$ between these two geodesics was at least π , their union would be a geodesic, contradicting the assumption. Thus we get that the angle at $q_{\hat{H}}$ is less than π . A short argument using [Lemma 8.15](#) then shows that the alternative projections are close. $\square$

Corollary 8.16. *For all $\delta > 0$ there exists $Q > 0$ such that this alternative projection is a Q -coarsely Lipschitz map from the set of points of $\partial\hat{T}_{\hat{c}}$ at least δ away from $\hat{H}$, with its intrinsic path metric, to $N^1(\hat{H})^{\text{sw}}$.*

8.6. A modification of the naive projection map. We wish to modify the naive projection to $N^1(\hat{H})^{\text{sw}}$ when the closest point projection to $\hat{H}$ lands in a horoball $\hat{T}_{\hat{c}}$. We will do this by using the alternative projection applied to the entry point to the horoball. The following helps us to understand the entry point.

Lemma 8.17. *There exists a universal constant C such that the following holds. Let x, y be points of $\hat{B}$, and let $\hat{T}_{\hat{c}}$ be a horoball. Suppose $x \notin \hat{T}_{\hat{c}}$ and $y \in \hat{T}_{\hat{c}}$. Let z be the point where the geodesic from x to y enters $\hat{T}_{\hat{c}}$. Then the distance from z to the closest point projection of x to $\hat{T}_{\hat{c}}$ is at most C .*

For a proof, compare to [PP10, Lemma 2.3], and note also that more precise bounds are possible.

Lemma 8.18. *For all $R > 0$ there exists an $R' > 0$ such that for any $\hat{T}_{\hat{c}}$, if two points of $\partial\hat{T}_{\hat{c}}$ are within distance R using the $\hat{B}$ metric, then they are within distance R' using the intrinsic path metric on $\partial\hat{T}_{\hat{c}}$.*

We leave it to the reader to prove Lemma 8.18 by starting with a geodesic in $\hat{B}$ and pushing it radially outwards from $\hat{c}$ until it gives a path on $\partial\hat{T}_{\hat{c}}$.

We can now give our definition of

$$\pi_{N^1(\hat{H})} : \hat{B}^\theta \rightarrow N^1(\hat{H})^{\text{sw}}.$$

Fix an ϵ large depending on Lemma 8.17 and Lemma 8.18 in ways that will become clear shortly. Consider $x \in \hat{B}^\theta$. If the closest point projection of x to $\hat{H}$ is not in a horoball, we define $\pi_{N^1(\hat{H})}(x)$ to be the naive projection. Similarly, if the projection is in a horoball, and the entry-point is within ϵ of $\hat{H}$, we define $\pi_{N^1(\hat{H})}(x)$ to be the naive projection. Otherwise, we look at the entry point z where the geodesic from x to $p_{\hat{H}}(x)$ enters $\hat{T}_{\hat{c}}$. We then define $\pi_{N^1(\hat{H})}(x)$ to be the alternative projection of z .

Remark 8.19. Continuing Remark 4.3, we note that the true domain $\hat{B}^{N^1(\hat{H})}$ of this map is the complement in $\hat{B}$ of the open r_θ neighborhood of $\hat{H}$ union all the interiors of all horoballs that $\hat{H}$ enters. (Here r_θ could be replaced with anything smaller, but some positive constant is required for Lemma 8.20. As a warning, we emphasize that $\hat{B}^{N^1(\hat{H})}$ is not convex, and as a rule arguments with angle domains require extra care.)

It remains to show the following.

Lemma 8.20. $\pi_{N^1(\hat{H})} : \hat{B}^{N^1(\hat{H})} \rightarrow \mathcal{C}(N^1(\hat{H}))$ is coarsely Lipschitz.

Proof. Consider points x, y that can be joined by a path in $\hat{B}^{N^1(\hat{H})}$ of length at most 1; it suffices to give a uniform bound on the distance between their images.

If the closest point projections $p_{\hat{H}}(x)$ and $p_{\hat{H}}(y)$ are not in a horoball, this follows from Lemma 8.11. If one of the closest point projections is not in a horoball and the other is in a horoball, it again follows from Lemma 8.11, given that the relevant entry point lies close to $\hat{H}$.

So assume that both of $p_{\hat{H}}(x)$ and $p_{\hat{H}}(y)$ are in a horoball $\hat{T}_{\hat{c}}$. Note that $p_{\hat{T}_{\hat{c}}}(x)$ and $p_{\hat{T}_{\hat{c}}}(y)$ are bounded distance from each other, and Lemma 8.17 gives that these points are bounded distance from associated entry points.

If these points are reasonably close to $\hat{H}$, the result follows from [Lemma 8.12](#) and [Lemma 8.11](#). If these points are not close to $\hat{H}$ it follows from [Corollary 8.16](#) and [Lemma 8.18](#). $\square$

8.7. Results for later use.

Lemma 8.21. *For any $\hat{V}$ as in [Section 5.2](#) and any $C > 0$ there exists a $D > 0$ such that the following holds. Suppose $y, z \in \hat{B}^\theta \cap \hat{V}$. Suppose also that $d_U(\pi_U(y), \pi_U(z)) < C$ for all angle domains U associated to hyperplanes passing through $\hat{V}$. Then the distance between y and z in the intrinsic $\hat{B}^\theta$ metric is at most D .*

Proof. On an appropriate pre-compact fundamental domain for the $\mathbb{Z}^k$ action on the set of points in $\hat{V}$ at least r_θ away from each hyperplane, the π_U are uniformly coarsely constant by [Lemma 8.20](#). The result now follows from [Lemma 8.10](#) and a compactness argument. $\square$

The following is an immediate consequence of [Lemma 8.12](#) and the definition of π_U .

Corollary 8.22. *Let U be an angle domain, let $\hat{H}$ be the corresponding hyperplane, and suppose $\hat{H}$ enters a horoball $\hat{T}_\varepsilon$. For any $x \in \partial\hat{T}_\varepsilon \cap \hat{B}^U$, the alternative projection of x to $\mathcal{C}(U)$ is uniformly close to $\pi_U(x)$.*

For later use we also record the following consequence of [Proposition 8.14](#) and the definition of $\pi_{N^1(\hat{H})}$.

Corollary 8.23. *There exists $C, D_0 > 0$ such that for any horoball $\hat{T}_\varepsilon$ and any $x, y \in \partial\hat{T}_\varepsilon$, the number of hyperplanes $\hat{H}$ intersecting $\hat{T}_\varepsilon$ whose angle domain U does not have $d_U(\pi_U(x), \pi_U(y)) < D_0$ is at most $Cd_{\partial\hat{T}_\varepsilon}(x, y) + C$, where $d_{\partial\hat{T}_\varepsilon}$ is the intrinsic path metric.*

We define “does not have $d_U(\pi_U(x), \pi_U(y)) < D_0$ ” to include the possibility that $\pi_U(x)$ or $\pi_U(y)$ are not defined, even when we extend the definition of π_U to all of $\partial\hat{T}_\varepsilon - \hat{H}$ in the natural way using the alternative projection.

Proof. We claim that $d_U(\pi_U(x), \pi_U(y)) < D_0$ for all angle domains U except those corresponding to singular points on the geodesic joining the images of x and y in one of the $\hat{\mathbb{C}}_i$. Indeed, suppose U is such an angle domain, corresponding to a hyperplane $\hat{H}$. By [Corollary 8.22](#), it suffices to prove the result for the alternative projection instead of π_U , so [Proposition 8.14](#) gives the claim.

The result then follows by uniform separation of the singular points in $\hat{\mathbb{C}}_i$. $\square$

9. NESTING AND ORTHOGONALITY

We now have a set of domains $\mathfrak{S}^b$ consisting of stratum domains, angle domains, and (center and tree) cusp domains, as well as associated maps $\pi_U : \hat{B}^\theta \rightarrow \mathcal{CU}$ to Gromov hyperbolic spaces. Much of the work so far can be summarized using [Definition 3.2](#) by noting that we have proven the following.

Proposition 9.1. *$(\hat{B}^\theta, \mathfrak{S}^b, \pi_\bullet)$ is a pre-HHS.*

Proof. For stratum domains, hyperbolicity was checked in [Lemma 5.7](#) and coarse Lipschitzness in the discussion immediately before. For center domains, hyperbolicity was checked in [Theorem 7.9](#), and coarse Lipschitzness at the start of [Section 7.5](#).

For tree domains, hyperbolicity was checked in [Lemma 7.5](#), and coarse Lipschitzness in the discussion immediately before. For angle domains, hyperbolicity was checked in [Corollary 8.8](#), and coarse Lipschitzness in [Lemma 8.20](#). $\square$

In this section we define the nesting and orthogonality relations and prove that $\mathfrak{S}^b$ is a nearly valid index set. Before proceeding, it is helpful to recall the notation for the different types of domains in $\mathfrak{S}^b$, which we have set up to allow more concise descriptions of these relations.

Stratum domains are associated to a stratum $\hat{S}^\circ$ of $\hat{B}$. These strata are typically not closed subsets of $\hat{B}$. We do not define stratum domains for 0-dimensional strata.

Angle domains are associated to the unit normal bundle $N^1(\hat{H})$ of a hyperplane $\hat{H}$.

Center domains are associated to horoballs $\hat{T}_c$.

Tree domains are associated to factors of the boundary of the horoball modulo the center direction. For each such factor, we formally define the associated domain to be the set of all hyperplanes that enter the horoball and are constant in that factor.

Warning 9.2. It is crucial to keep in mind that the terms “nested” and “orthogonal” describe the map $\pi_U \times \pi_V$ and not geometric objects in $\hat{B}$ associated to U and V . See [Warning 9.5](#) for examples of how otherwise these terms can cause confusion.

9.1. The definition of nesting and orthogonality. The nesting relation is defined as follows.

- (1) A stratum domain is nested in a second stratum domain if the first stratum is contained in the closure of the second.
- (2) An angle domain is nested in a stratum domain if the closure of the stratum intersects but is not contained in the hyperplane associated to the angle domain.
- (3) A center domain is nested in a stratum domain if the stratum enters the associated horoball.
- (4) A tree domain is nested in a stratum domain if the stratum enters the associated horoball and is not constant in the associated factor.
- (5) An angle domain is nested in a tree domain if the hyperplane associated to the angle domain enters the horoball associated to the tree domain and is constant in the factor associated to the tree domain.

Remark 9.3. Angle domains and center domains are always nest-minimal, meaning they have nothing properly nested in them. Only angle domains can be nested in tree domains. Tree domains are never nest minimal.

Remark 9.4. Strata are in general not closed and are in general not cobounded. But strata can in exceptional cases be cocompact, and a stratum domain is nest minimal if and only if the stratum is cocompact. (Keep in mind that 0-dimensional strata do not give domains, so in particular do not give nest-minimal domains.)

The orthogonality relation is defined as follows.

- (1) A stratum domain is orthogonal to an angle domain if the stratum is contained in the associated hyperplane.

- (2) Distinct angle domains are orthogonal if their associated hyperplanes intersect.
- (3) An angle domain is orthogonal to a center domain if the associated hyperplane enters the associated horoball.
- (4) An angle domain is orthogonal to a tree domain if the associated hyperplane enters the associated horoball and is not constant in the associated factor.
- (5) A center domain is orthogonal to a tree domain if they are both associated to the same horoball.
- (6) Distinct tree domains are orthogonal if they are both associated to the same horoball.

As always, two domains are asynchronous if neither is nested in the other and they are not orthogonal. The relations are summarized in Figure 8, where a blank entry indicates the domains cannot satisfy the relation.

$\text{top} \perp \text{side}$ $\text{top} \sqsubset \text{side}$ $\text{top} \triangleleft \text{side}$	stratum domain $\hat{S}^\circ$	angle domain $N^1(\hat{H}')$	center domain $\hat{T}_{e'}$	tree domain V' {hyperplanes constant in fixed factor of $\hat{T}_{e'}$ }
stratum domain $\hat{S}^\circ$	$\hat{S}' \subset \hat{S}$ $\hat{S}' \not\subset \hat{S} \ \& \ \hat{S} \not\subset \hat{S}'$	$\hat{S} \subset \hat{H}'$ $\hat{S} \not\subset \hat{H}' \ \& \ \hat{S} \cap \hat{H}' \neq \emptyset$ $\hat{S} \cap \hat{H}' = \emptyset$	$\hat{S} \cap \hat{T}_{e'} \neq \emptyset$ $\hat{S} \cap \hat{T}_{e'} = \emptyset$	$\hat{S} \cap \hat{T}_{e'} \neq \emptyset \ \& \ \hat{S} \not\subset \bigcup V'$ $\hat{S} \cap \hat{T}_{e'} = \emptyset \ \text{or} \ \hat{S} \subset \bigcup V'$
angle domain $N^1(\hat{H})$	$\hat{S}' \subset \hat{H}$ $\hat{H} \cap \hat{S}' = \emptyset$	$\hat{H} \cap \hat{H}' \neq \emptyset \ \& \ \hat{H} \neq \hat{H}'$ $\hat{H} \cap \hat{H}' = \emptyset$	$\hat{H} \cap \hat{T}_{e'} \neq \emptyset$ $\hat{H} \cap \hat{T}_{e'} = \emptyset$	$\hat{H} \cap \hat{T}_{e'} \neq \emptyset \ \& \ \hat{H} \notin V'$ $\hat{H} \cap \hat{T}_{e'} = \emptyset$
center domain $\hat{T}_e$	$\hat{S}' \cap \hat{T}_e = \emptyset$	$\hat{H}' \cap \hat{T}_e \neq \emptyset$ $\hat{H}' \cap \hat{T}_e = \emptyset$	$\hat{e} \neq \hat{e}'$	$\hat{e} = \hat{e}'$ $\hat{e} \neq \hat{e}'$
tree domain V {hyperplanes constant in fixed factor of $\hat{T}_e$ }	$\hat{S}' \cap \hat{T}_e = \emptyset \ \text{or} \ \hat{S}' \subset \bigcup V$	$\hat{H}' \cap \hat{T}_e \neq \emptyset \ \& \ \hat{H}' \notin V$ $\hat{H}' \in V$ $\hat{H}' \cap \hat{T}_e = \emptyset$	$\hat{e} = \hat{e}'$ $\hat{e} \neq \hat{e}'$	$\hat{e} = \hat{e}' \ \& \ V \neq V'$ $\hat{e} \neq \hat{e}'$

FIGURE 8. The definition of orthogonal, properly nested, and asynchronous.

Warning 9.5. We now repeat [Warning 9.2](#) in more detail to emphasize that one may not think of the words “nesting” and “orthogonality” as naively describing the geometric objects associated to the domains. For example, nesting doesn’t always correspond to containment and orthogonality doesn’t always correspond to perpendicular intersections. Note in particular that an angle domain defined by a hyperplane is orthogonal to a stratum domain when the hyperplane contains the stratum.

9.2. First observations on the relations.

Lemma 9.6. *Every chain $U_1 \sqsubset U_2 \sqsubset \dots \sqsubset U_\ell$ that isn’t properly contained in a larger chain has the following properties.*

- (1) $U_\ell = \hat{B}^\circ$.
- (2) The $U_i, i \geq 3$ are stratum domains.
- (3) U_2 is a stratum domain or a tree domain.
- (4) If U_2 is a tree domain then U_1 is an angle domain.

Proof. The first claim is true because every domain is nested in $\widehat{B}^\circ$. The other claims follow from [Remark 9.3](#). $\square$

Corollary 9.7. *Nesting is a partial order.*

Proof. Suppose $U \subsetneq V$ and $V \subsetneq W$. We need to show $U \subsetneq W$.

Note that W must be a stratum domain, and V must be a stratum domain or a tree domain. We leave the case where V is a stratum domain to the reader and do the case when V is a tree domain. In this case U is an angle domain. We know

- (1) the hyperplane associated to U enters the horoball associated to V and is constant in the factor associated to V , and
- (2) the stratum associated to W enters the horoball associated to V and is not constant in the factor associated to V .

Inside of a horoball, hyperplanes are defined by setting a single coordinate equal to a constant, and the closures of strata are defined by setting some number of coordinates equal to a constant. Thus we can conclude that the stratum closure $\overline{W}$ and the hyperplane associated to U intersect and the former is not contained in the latter. Hence $U \subsetneq W$ as desired. $\square$

Corollary 9.8. *The maximal size ℓ of a chain $U_1 \subsetneq U_2 \subsetneq \cdots \subsetneq U_\ell$ is at most $n+1$, and is equal to $n+1$ if and only if Γ is non-uniform or there is a 0-dimensional stratum.*

Proof. In the non-uniform case, [Assumption \(D\)](#) gives that there exists a 1-dimensional stratum U_1 that enters a horoball. We can thus find a center domain U_0 and stratum domains $U_1, \dots, U_n$ so

$$U_0 \subsetneq U_1 \subsetneq \cdots \subsetneq U_n,$$

proving $\ell \geq n+1$.

If there is a 0-dimensional stratum, it is defined as the intersection of n distinct hyperplanes, say $\widehat{H}_1, \widehat{H}_2, \dots, \widehat{H}_n$. For $i = 1, \dots, n$ we can let U_i be the generic stratum of $\cap_{j=1}^{n-i} \widehat{H}_j$, and we can let U_0 be the angle domain $N^1(\widehat{H}_n)$, again proving $\ell \geq n+1$.

The remaining claims follow from [Lemma 9.6](#). $\square$

The following statement is slightly complicated by our choice not to define stratum domains for 0-dimensional strata.

Lemma 9.9. *A set of domains is an orthogonal set not contained in a larger orthogonal set if and only if it is one of the following.*

- (1) A collection of n angle domains corresponding to n distinct hyperplanes through a point.
- (2) A stratum domain and the collection of angle domains corresponding to hyperplanes containing it.
- (3) The n cusp domains associated to a horoball, as well as the same with any subset of the tree domains each replaced by an angle domain nested in it.

Proof. A stratum domain is only orthogonal to angle domains corresponding to hyperplanes containing the stratum, and all these angle domains happen to be orthogonal to each other.

Similarly a center domain is only orthogonal to the tree domains associated to the same horoball and the angle domains nested in these tree domains. Again these

tree domains happen to be orthogonal to each other, and the same is true if some of them are replaced with an angle domain nested in them.

This proves the claim if the tuple contains a stratum domain or a center domain. Anything orthogonal to a tree domain is orthogonal to the associated center domain, so this also proves it unless the tuple consists entirely of angle domains. Any orthogonal set of angle domains must correspond to hyperplanes with a common intersection, by [Lemma 2.6](#) and [Remark 5.10](#), and one can add the generic stratum of this intersection to the orthogonal set as long as this intersection is not a point. $\square$

Corollary 9.10. *The maximum size of an orthogonal set of domains is at most n . It is equal to n if and only if Γ is non-uniform or there is a 0-dimensional stratum or a 1-dimensional stratum or $n = 1$.*

Remark 9.11. Suppose $n > 1$. Then the only domain not orthogonal to any other is the nest maximal domain.

9.3. Nearly valid. We require two more lemmas before we can conclude that $\mathfrak{S}^b$ is a nearly valid index set.

Lemma 9.12. *If $V \subseteq W$ and $W \perp U$ then $V \perp U$.*

Proof. It suffices to assume $V \neq W$. Since W has something properly nested in it, it must be a stratum domain or a tree domain.

Case 1: W is a stratum domain. Then U is an angle domain, and we leave it to the reader to check the result for each of the four possible types of V .

Case 2: W is a tree domain. In this case U is a center domain, or a tree domain corresponding to one of the other factors of the horoball, or an angle domain nested in one of those other tree domains, and V is an angle domain. Again we leave the details to the reader. $\square$

Lemma 9.13. *If $V \subseteq W$ then the set $\mathcal{P}_{V,W}$ of domains orthogonal to V and nested in W has at most finitely many nest maximal elements, all of which are orthogonal to each other and to V .*

Proof. We first prove this when $W = \hat{B}^\circ$ is the nest maximal element. In this case:

- (1) If V is a stratum domain, $\mathcal{P}_{V,W}$ will be an orthogonal tuple of angle domains.
- (2) If V is a cusp domain, the elements of $\mathcal{P}_{V,W}$ are the $n - 1$ other cusp domains and, for those that are tree domains, the angle domains nested in them. So the maximal elements of $\mathcal{P}_{V,W}$ together with V itself will be the collection of n orthogonal cusp domains associated to a single horoball.
- (3) If V is an angle domain, then V is orthogonal only to the generic stratum of the hyperplane V is associated to and the domains nested in that stratum. So $\mathcal{P}_{V,W}$ has only a single nest maximal element, and it is orthogonal to V .

More generally, we must consider the case when W is an arbitrary stratum domain and when it is a tree domain. We leave the stratum case to the reader. The case when W is a tree domain doesn't really occur, since only angle domains can be nested in a tree domain, and two angle domains nested in the same tree domain are never orthogonal: so $\mathcal{P}_{V,W}$ is always empty in this case. $\square$

We now have near validity of the index set.

Proposition 9.14. $\mathfrak{S}^b$ with the relations defined above is a nearly valid index set.

Proof. Orthogonality is symmetric and anti-reflexive by definition; nesting is a partial order by [Corollary 9.7](#), with a unique maximal element by [Lemma 9.6](#). The finite height axiom is true by [Corollary 9.8](#); and the coherence axiom by [Lemma 9.12](#). The weak containers axiom is true by [Lemma 9.13](#) and [Corollary 9.10](#). $\square$

10. ACTIVE REGIONS

In this section we establish the density, bounded projections, and structured redundancy assumptions of [Proposition 3.6](#).

We define the active regions $\mathcal{A}_U$ for [Proposition 3.6](#) as follows. To start, define closed convex subsets K_U of $\hat{B}$ as follows.

- (1) If U is a stratum domain, K_U is the associated stratum closure.
- (2) If U is an angle domain, K_U is the associated hyperplane.
- (3) If U is a cusp domain, K_U is the associated horoball.

We then define $\mathcal{B}_U$ to be the $10nr_\theta$ -neighborhood of K_U , and we define $\mathcal{A}_U = \mathcal{B}_U \cap \hat{B}^\theta$. We also recall from [Remarks 4.3, 5.6, 7.4, 7.6](#) and [8.19](#) that $\hat{B}^U$ satisfies $\hat{B}^\theta \subset \hat{B}^U \subset \hat{B}$ and is the domain of π_U .

10.1. Density. We start by observing the following.

Lemma 10.1. *The density assumption of [Proposition 3.6](#) holds.*

Proof. It suffices to check this for maximal orthogonal sets, which [Lemma 9.9](#) gives are one of the following.

- (1) A collection of n angle domains corresponding to n distinct hyperplanes through a point.
- (2) A stratum domain and the collection of angle domains of codimension 1 strata containing it.
- (3) The n cusp domains associated to a horoball, as well as the same with any subset of the tree domains each replaced by an angle domain nested in it.

In the first case, use [Remark 5.2](#) to find a thick point near the point of intersection of the n -hyperplanes and then apply [Lemma 8.10](#). The other cases follow immediately by additionally applying the discussion in [Sections 5](#) and [7](#). $\square$

10.2. Bounded projections. We now address the bounded projection assumption. We actually prove something slightly stronger: Rather than using $\mathcal{A}_U = \mathcal{B}_U \cap \hat{B}^\theta$, we use the larger set $\mathcal{B}_U \cap \hat{B}^V$.

Proposition 10.2. *If $U \triangleleft V$ or $U \sqsubset V$, $\text{diam } \pi_V(\mathcal{B}_U \cap \hat{B}^V)$ is uniformly bounded.*

To prove [Proposition 10.2](#), we need the following, which combines a restatement of [Lemma 2.7](#) and an immediate consequence of [Lemma 2.7](#).

Corollary 10.3. *For all $\epsilon > 0$ there exists a $\delta > 0$ such that the following holds: Suppose C is a closed convex subset of a $\text{CAT}(-1)$ space X and p_C is projection to C . Suppose A is a convex set.*

- (1) *If A is disjoint from the ϵ -neighborhood of C , then $p_C(A)$ has diameter at most δ , both in the ambient metric and in the induced path metric on the boundary of C .*

- (2) If A is not disjoint from the ϵ -neighborhood of C , then $p_C(A)$ is contained in the δ -neighborhood of the intersection of A and the ϵ -neighborhood of C .

Throughout this section, angle domains need to be handled with special care, due to the modified nature of their projections. So we start with the following.

Lemma 10.4. *For all $\epsilon > 0$ there exists $C > 0$ such that the following holds. Suppose $x, y \in \hat{B}^U$, where U is an angle domain with corresponding hyperplane $\hat{H}$, and the geodesic from x to y does not come within distance ϵ of $\hat{H}$ and does not enter any horoball that $\hat{H}$ enters. Then $d_U(\pi_U(x), \pi_U(y)) \leq C$.*

Proof. First suppose $\pi_U(x)$ and $\pi_U(y)$ are equal to the naive projections; so if the nearest point projections of x and y lie in a horoball that $\hat{H}$ enters, the entry point is reasonably close to $\hat{H}$. Then Lemma 8.11 gives the result.

Suppose the nearest point projection of x to $\hat{H}$ lies in a horoball $\hat{T}_{\hat{e}}$ that $\hat{H}$ enters. Corollary 10.3 gives that the nearest point projections of x and y to $\hat{H}$ are close to each other, and a small elaboration of Corollary 10.3 gives the same for the projections to $\hat{T}_{\hat{e}}$. If these closest point projections to $\hat{T}_{\hat{e}}$ are far from $\hat{H}$, the result follows from the fact that π_U is coarsely Lipschitz (Lemma 8.20), because Lemma 8.17 gives that the entry points are close to each other.

Lemma 8.12 gives that if $z \in \partial \hat{T}_{\hat{e}}$ has bounded distance from $\hat{H}$ then the naive and modified projections of z are bounded distance apart. Thus, in the remaining case it suffices to use the naive projection, and again Lemma 8.11 gives the result. $\square$

Proof of Proposition 10.2. For convenience, we reproduce a modified form of Figure 8 here as Figure 9, which indicates which cases must be handled. We consider

top $\sqsubset$ side top $\triangleright$ side	stratum domain $\hat{S}'$	angle domain $N^1(\hat{H}')$	center domain $\hat{T}_{\hat{e}'}$	tree domain W' {hyperplanes constant in fixed factor of $\hat{T}_{\hat{e}'}$ }
stratum domain $\hat{S}^\circ$	$\hat{S}' \subseteq \hat{S}$ $\hat{S}' \not\subseteq \hat{S} \text{ \& } \hat{S} \not\subseteq \hat{S}'$	$\hat{S} \not\subseteq \hat{H}' \text{ \& } \hat{S} \cap \hat{H}' \neq \emptyset$ $\hat{S} \cap \hat{H}' = \emptyset$	$\hat{S} \cap \hat{T}_{\hat{e}'} \neq \emptyset$ $\hat{S} \cap \hat{T}_{\hat{e}'} = \emptyset$	$\hat{S} \cap \hat{T}_{\hat{e}'} \neq \emptyset \text{ \& } \hat{S} \not\subseteq \bigcup W'$ $\hat{S} \cap \hat{T}_{\hat{e}'} = \emptyset \text{ or } \hat{S} \subseteq \bigcup W'$
angle domain $N^1(\hat{H})$	$\hat{H} \cap \hat{S}' = \emptyset$	$\hat{H} \cap \hat{H}' = \emptyset$	$\hat{H} \cap \hat{T}_{\hat{e}'} = \emptyset$	$\hat{H} \cap \hat{T}_{\hat{e}'} = \emptyset$
center domain $\hat{T}_{\hat{e}}$	$\hat{S}' \cap \hat{T}_{\hat{e}} = \emptyset$	$\hat{H}' \cap \hat{T}_{\hat{e}} = \emptyset$	$\hat{e} \neq \hat{e}'$	$\hat{e} \neq \hat{e}'$
tree domain W {hyperplanes constant in fixed factor of $\hat{T}_{\hat{e}}$ }	$\hat{S}' \cap \hat{T}_{\hat{e}} = \emptyset \text{ or } \hat{S}' \subseteq \bigcup W$	$\hat{H}' \in W$ $\hat{H}' \cap \hat{T}_{\hat{e}} = \emptyset$	$\hat{e} \neq \hat{e}'$	$\hat{e} \neq \hat{e}'$

FIGURE 9. The definition of properly nested and asynchronous.

cases where V is a domain on the side and U is a domain on the top and the conditions given in Figure 9 hold, and we must show $\pi_V(\mathcal{B}_U \cap \hat{B}^V)$ is bounded. We do this in four cases, each of which starts off by establishing this claim.

Claim: Every point of $\mathcal{B}_U \cap \hat{B}^V$ has uniformly bounded distance in the intrinsic $\hat{B}^V$ path metric to a point of $K_U \cap \hat{B}^V$.

Case 1: V is a stratum domain. Let $\hat{S}$ be the associated stratum closure. In this case $\hat{B}^V = \hat{B}$, and K_U is one of: a horoball; or a stratum closure not containing

$\hat{S}$. (Hyperplanes are stratum closures.) Since $\hat{B}^V = \hat{B}$ the claim is immediate in this case.

By the claim, it suffices to show that $\pi_V(K_U)$ is uniformly bounded. Here π_V is the composition of closest point projection to $\hat{S}$ followed by inclusion into the electrification. [Corollary 10.3](#) and [Remarks 5.9](#) and [5.10](#) give that the closest point projection of K_U to $\hat{S}$ is either uniformly bounded or uniformly close to being contained in a smaller stratum closure or a horoball intersected with $\hat{S}$. Since these regions of $\hat{S}$ are contained in regions that are electrified in $\mathcal{C}(V)$, it follows that $\pi_V(K_U)$ is uniformly bounded.

Case 2: V is a center domain. Let $\hat{T}_\varepsilon$ be the associated horoball. In this case $\hat{B}^V = \hat{B} - \text{int}(\hat{T}_\varepsilon)$, and K_U is one of: a stratum closure disjoint from $\hat{T}_\varepsilon$; or a horoball disjoint from $\hat{T}_\varepsilon$. The claim is immediate, since [Remark 4.2](#) provides more than enough separation from $\hat{T}_\varepsilon$.

Keeping in mind [Remark 4.2](#), [Corollary 10.3](#) gives that the closest point projection of K_U to $\hat{T}_\varepsilon$ is bounded. This gives the result because π_V factors through that projection and π_V is coarsely Lipschitz.

Case 3: V is a tree domain. Let $\hat{T}_\varepsilon$ be the associated horoball. In this case π_V is defined on the complement of the interior of $\hat{T}_\varepsilon$, and K_U is one of: a horoball disjoint from $\hat{T}_\varepsilon$; a stratum closure disjoint from $\hat{T}_\varepsilon$; or a stratum closure which is constant in the factor $\hat{\mathbb{C}}_i$ associated to V .

For both the claim and the result, the first two of the possibilities for K_U follow as in Case 2, so let us consider the case of a stratum closure $\hat{S}'$ which is constant in $\hat{\mathbb{C}}_i$. For the claim, this last case is not much harder, except that for points of $\mathcal{B}_U \cap \hat{B}^V$ very close to $\hat{T}_\varepsilon$ the geodesic to K_U , which is very short, may require a slight radial push outwards if it enters $\hat{T}_\varepsilon$.

By [Remark 5.9](#), the closest point projection to $\hat{T}_\varepsilon$ of $\hat{S}' - \text{int}(\hat{T}_\varepsilon)$ is equal to $\hat{S}' \cap \partial\hat{T}_\varepsilon$. Since $\hat{S}' \cap \partial\hat{T}_\varepsilon$ is constant in the relevant factor $\hat{\mathbb{C}}_i$ the result follows.

Case 4: V is an angle domain. Let $\hat{H}$ be the associated hyperplane. In this case π_V is defined on the set of points at least r_θ from $\hat{H}$ and not in the interior of a horoball that $\hat{H}$ enters, and K_U is one of: a stratum closure disjoint from $\hat{H}$; or a horoball that is disjoint from $\hat{H}$.

For the claim, the key subcase here is when K_U is a stratum closure that enters a horoball that $\hat{H}$ enters, and this subcase follows as in Case 3 since [Remark 4.2](#) gives a separation of at least $1000nr_\theta$ between $\hat{H}$ and the stratum closure outside of horoballs. The other subcases follow as in Case 2. This proves the claim in this case.

To prove the result, we divide into two subcases.

Case 4a: No horoball is entered by both K_U and $\hat{H}$. In this case we get the result by [Lemma 10.4](#) and [Remark 4.2](#).

Case 4b: Some horoball is entered by both K_U and $\hat{H}$. In this case, K_U is a stratum closure $\hat{S}$ that is disjoint from $\hat{H}$ but enters a horoball $\hat{T}_\varepsilon$ that $\hat{H}$ enters. Note that $\hat{S}$ cannot go into any other horoball that $\hat{H}$ enters, since otherwise $\hat{S}$ and $\hat{H}$ would have a geodesic in common [[DSU17](#), Proposition 4.4.4].

Every point of $K_U - \hat{T}_\varepsilon$ can be joined to its closest point projection to $\hat{T}_\varepsilon$ by a geodesic. This closest point projection lies in $\hat{S} \cap \partial\hat{T}_\varepsilon$ by [Remark 5.9](#), so the geodesic

lies in the region where Lemma 10.4 applies. The intersection $\hat{S} \cap \partial \hat{T}_{\hat{e}}$ maps to a coarse point of $\mathcal{C}(V)$ via the alternative projection, by Lemma 8.15. This gives the result by Corollary 8.22. $\square$

We conclude with the auxiliary statement of the boundedness assumption.

Lemma 10.5. *If $U \subsetneq V$ then $\mathcal{A}_U \cap \mathcal{A}_V \neq \emptyset$.*

Proof. This follows immediately from the definitions; see Figure 9. $\square$

10.3. Structured redundancy. We prove the two halves of the structured redundancy assumption separately.

Proposition 10.6. *There exists a $B > 0$ such that if $U \triangleleft V$, for any $x \in \hat{B}^\theta$ we have at least one of*

$$d_U(\pi_U(x), \pi_U(\mathcal{A}_V)) \leq B \quad \text{or} \quad d_V(\pi_V(x), \pi_V(\mathcal{A}_U)) \leq B.$$

Roughly speaking, the proof follows a common strategy that considers a sort of first entry point.

Proof. To prove this, start by defining $\mathcal{B}_U^+$ to be $\mathcal{B}_U$ if U is a stratum or cusp domain, and to be $\mathcal{B}_U$ union all the horoballs it intersects if U is an angle domain.

Let z be any point of $\mathcal{A}_U \cup \mathcal{A}_V$. Let y be the first point along the geodesic from x to z that is contained in $\mathcal{B}_U^+ \cup \mathcal{B}_V^+$. Without loss of generality, assume $y \in \mathcal{B}_U^+$. If y is in both $\mathcal{B}_U^+$ and $\mathcal{B}_V^+$, then without loss of generality assume that if y is in $\mathcal{B}_U \cup \mathcal{B}_V$ then $y \in \mathcal{B}_U$.

Note that the interior of the geodesic from x to y is disjoint from $\mathcal{B}_U^+ \cup \mathcal{B}_V^+$, and so in particular the geodesic is in $\hat{B}^U \cap \hat{B}^V$. By Proposition 10.2, it suffices to show $\pi_V(x)$ is close to $\pi_V(\mathcal{B}_U \cap \hat{B}^V)$ or $\pi_U(x)$ is close to $\pi_U(\mathcal{B}_V \cap \hat{B}^U)$.

Case 1: $y \in \mathcal{B}_U$. In this case $\pi_V(\mathcal{B}_U \cap \hat{B}^V) \approx \pi_V(x)$ follows from Corollary 10.3 and Lemma 10.4.

Note: If Case 1 does not apply, then since $y \notin \mathcal{B}_U \cup \mathcal{B}_V$ and $y \in \mathcal{B}_U^+$ we have

- (1) U is an angle domain and the corresponding hyperplane enters a horoball $\hat{T}_{\hat{e}}$ such that $y \in \partial \hat{T}_{\hat{e}}$, and
- (2) V cannot be a cusp domain associated to $\hat{T}_{\hat{e}}$.

We assume this in all the remaining cases.

Case 2: V is a cusp domain, or V is a stratum or angle domain with K_V disjoint from $\hat{T}_{\hat{e}}$. In this case Remark 4.2 gives that the set K_V is far from $\hat{T}_{\hat{e}}$. Since there are points of $\mathcal{B}_U \cap \hat{B}^V$ arbitrarily close to $\hat{T}_{\hat{e}}$, $\pi_V(\mathcal{B}_U \cap \hat{B}^V) \approx \pi_V(x)$ again follows from Corollary 10.3, Lemma 10.4, and Remark 4.2.

Case 3: V is a stratum domain whose associated stratum intersects $\hat{T}_{\hat{e}}$. Since there are points of $\mathcal{B}_U \cap \hat{B}^V$ arbitrarily close to $\hat{T}_{\hat{e}}$, and since $\pi_V(\hat{T}_{\hat{e}})$ is a coarse point, we get $\pi_V(\mathcal{B}_U \cap \hat{B}^V) \approx \pi_V(x)$ from Corollary 10.3.

Case 4: V is an angle domain whose hyperplane intersects $\hat{T}_{\hat{e}}$. So the two hyperplanes, which must be disjoint, both enter the given horoball. We continue to make use of the entry point y defined above.

We do a second version of the rough idea used above, but now in the factor $\hat{\mathbb{C}}_i$ where the hyperplanes associated to U and V are constant. Let $q_U, q_V \in \hat{\mathbb{C}}_i$ be

those constant values, and let y_i be the image of y in $\hat{\mathbb{C}}_i$. Pick $z' \in \{q_U, q_V\}$, and consider the geodesic in $\hat{\mathbb{C}}_i$ from y_i to z' . Let y'_i be the first point of this geodesic in $\{q_U, q_V\}$, and without loss of generality assume $y'_i = q_U$.

[Lemma 10.4](#), [Proposition 8.14](#), and [Corollary 8.22](#) then give that $\pi_V(\mathcal{B}_U \cap \hat{B}^V) \approx \pi_V(y) \approx \pi_V(x)$. $\square$

Proposition 10.7. *There exists a $B > 0$ such that the following holds. If $U \subsetneq V$, for any $x, y \in \hat{B}^\theta$, if there is a geodesic from $\pi_V(x)$ to $\pi_V(y)$ that has distance at least B from $\pi_V(\mathcal{A}_U)$, then $d_U(\pi_U(x), \pi_U(y)) \leq B$.*

Proof. We divide the proof into cases.

Observation 1: If W is a stratum domain with corresponding stratum $\hat{S}^\circ$, π_W is defined on $\hat{B}$. For any x, y , the geodesic from $\pi_W(x)$ to $\pi_W(y)$ is Hausdorff close to π_W of the geodesic in $\hat{B}$ from x to y , by [Corollary 10.3](#) and [Proposition 2.3](#).

Case 1: V is a stratum domain and U is a stratum domain. In this case we have a stratum $\hat{S}^\circ$ associated to V and another stratum contained in the closure of $\hat{S}^\circ$, and it suffices to apply observation 1 and [Corollary 10.3](#).

Case 2: V is a stratum domain and U is a cusp domain. In this case we have a stratum $\hat{S}^\circ$ associated to V , and a horoball $\hat{T}_\epsilon$ that $\hat{S}^\circ$ enters (with an extra condition if U is a tree domain). Suppose that $d_U(\pi_U(x), \pi_U(y))$ is large. It follows that the closest point projections of x and y to $\hat{T}_\epsilon$ are far apart, and hence that the geodesic in $\hat{B}$ from x to y enters $\hat{T}_\epsilon$. Observation 1 now gives that there is a geodesic from $\pi_V(x)$ to $\pi_V(y)$ that passes close to $\pi_V(\mathcal{A}_U)$.

Case 3: V is a stratum domain and U is an angle domain. In this case we have a stratum $\hat{S}^\circ$ associated to V and a hyperplane $\hat{H}$ associated to U , and $\hat{S}$ intersects but is not contained in $\hat{H}$. Consider x and y whose images in $\mathcal{C}(U) = N^1(\hat{H})^{\text{sw}}$ are far apart.

First suppose that the $\hat{B}$ geodesic from x to y doesn't go into any horoball that $\hat{H}$ goes into. In this case [Lemma 10.4](#) shows that the geodesic from x to y must come close to $\hat{H}$. This gives that the geodesic from $\pi_V(x)$ to $\pi_V(y)$ must come close to $\pi_V(\mathcal{A}_U)$.

Next suppose that the geodesic from x to y does go into a horoball $\hat{T}_\epsilon$ that $\hat{H}$ goes into. If $\hat{S}$ also goes into $\hat{T}_\epsilon$, this directly gives that the geodesic from $\pi_V(x)$ to $\pi_V(y)$ must come close to $\pi_V(\mathcal{A}_U)$, since the horoball is electrified in $\mathcal{C}(V)$. If $\hat{S}$ does not go into $\hat{T}_\epsilon$, it suffices to note that the projection of $\hat{T}_\epsilon$ to $\hat{S}$ is bounded and is uniformly bounded distance from $\pi_V(\mathcal{A}_U)$.

Case 4: V is a tree domain and U is an angle domain. In this case the hyperplane $\hat{H}$ associated to U intersects the horoball $\hat{T}_\epsilon$ associated to V , and $\hat{H}$ is constant in the factor $\hat{\mathbb{C}}_i$ associated to V . Assume that $\pi_U(x)$ and $\pi_U(y)$ are far apart.

Case 4a: Suppose there is a geodesic from either x or y to $\hat{H}$ that doesn't go into $\hat{T}_\epsilon$. Without loss of generality, say there is a geodesic from x to a point x' of $\hat{H}$ that doesn't go into $\hat{T}_\epsilon$. We get that $\pi_V(x)$ is about $\pi_V(x')$, which is uniformly close to $\pi_V(\mathcal{A}_U)$.

Thus, in the remaining subcases, we can restrict to the cases where the closest point projections of x and y to $\hat{H}$ lie in $\hat{T}_\epsilon$, and a small elaboration of the argument

above also allows us to restrict to the case where these closest point projections of x and y to $\hat{H}$ are not close to the complement of $\hat{T}_{\hat{c}}$.

Case 4b: Suppose the $\hat{B}$ geodesic from x to y doesn't go into $\hat{T}_{\hat{c}}$. In this case, this geodesic must go close to $\hat{H}$ or a horoball that $\hat{H}$ enters other than $\hat{T}_{\hat{c}}$, by [Lemma 10.4](#). The projection of $\hat{H}$ or any such horoball is coarsely equal to $\pi_V(\mathcal{A}_U)$, so this gives that π_V of this geodesic must come close to $\pi_V(\mathcal{A}_U)$. Since we are assuming the geodesic doesn't go into $\hat{T}_{\hat{c}}$, this image under π_V is coarsely constant, thus giving in a degenerate form the desired conclusion that the geodesic from $\pi_V(x)$ to $\pi_V(y)$ must come close to $\pi_V(\mathcal{A}_U)$.

Case 4c: Suppose the $\hat{B}$ geodesic from x to y does go into $\hat{T}_{\hat{c}}$. If it enters or exits near $\hat{H}$, the result follows from [Lemma 8.17](#), so we can assume otherwise. By supposition, π_U is computed using the modified part of the definition, so [Proposition 8.14](#) gives the result. $\square$

11. PROOF OF HIERARCHICAL HYPERBOLICITY

In this section we prove the enough projections and progress localization assumptions of [Proposition 3.6](#), and then we put it all together to prove the main results of this paper.

11.1. Enough projections. We start with the enough projections assumption.

Proposition 11.1. *For each $\kappa \geq 0$ there exists $\omega \geq 0$ such that if $x, y \in \hat{B}^\theta$ satisfy $d_{\hat{B}^\theta}(x, y) \geq \omega$ then there exists U such that*

$$d_U(\pi_U(x), \pi_U(y)) \geq \kappa.$$

We isolate the largest part of the remaining required analysis as a lemma.

Lemma 11.2. *For each $\kappa' \geq 0$ there exists $\omega' \geq 0$ such that if $x, y \in \hat{B}^\theta$ satisfy $d_{\hat{B}}(x, y) \leq \kappa'$ and $d_U(\pi_U(x), \pi_U(y)) \leq \kappa'$ for all angle domains U , then $d_{\hat{B}^\theta}(x, y) \leq \omega'$.*

Proof. For expositional clarity we start with the case when Γ is cocompact, or equivalently when there are no cusp domains.

Throughout this proof, we will refer to $\hat{V}$ as in [Section 5.2](#), and we will always assume this $\hat{V}$ belongs to a finite union of $\hat{\Gamma}$ orbits such that the $\hat{V}_0$ associated to this finite union of $\hat{\Gamma}$ orbits covers $\hat{B}$. (More generally they should cover the complement of the horoballs, but for the moment we assume there aren't any horoballs.)

Consider the geodesic in $\hat{B}$ from x to y . Pick a sequence

$$x = x_0, x_1, \dots, x_\ell = y$$

of $O(d_{\hat{B}}(x, y))$ points along this geodesic, in order, each distance at most $\xi/100$ from the next. For each i , find one of the $\hat{V}_0$ from [Section 5.2](#) that contains x_i , call it $\hat{V}_0^{(i)}$. Recall these $\hat{V}_0$ are balls of radius $r_i \geq \xi/100$ about some center point p_i contained in every hyperplane that intersects $\hat{V}$, and that the ball of radius $100r_i$ about p_i is contained in $\hat{V}$.

Using [Remark 5.2](#) on the point p_i , find a point x'_i whose distance to p_i is much less than $\xi/100$. Using [Lemma 8.10](#), adjust it to a point $x''_i \in \hat{B}^\theta$ with $d(p_i, x''_i) = d(p_i, x'_i)$ and such that $\pi_U(x''_i)$ is uniformly close to $\pi_U(x)$ for all angle domains U

associated to hyperplanes passing through $\widehat{V}$. The fact that π_U is coarsely Lipschitz applied either to the geodesic $[x, x_i]$ or the geodesic $[x_i, y]$, implies that same for all angle domains U associated to hyperplanes not passing through $\widehat{V}$.

We can assume $x''_0 = x$ and $x''_\ell = y$. The triangle inequality gives

$$d_{\widehat{B}}(p_i, x''_{i+1}) \leq d(p_i, x_i) + d(x_i, x_{i+1}) + d(x_{i+1}, x''_{i+1}) \leq 2r_i + \xi/100 + 2r_{i+1}.$$

We also get the same for $d_{\widehat{B}}(p_{i+1}, x''_i)$, so depending on which of r_i, r_{i+1} is larger we get that either both x''_i and x''_{i+1} are in $\widehat{V}^{(i)}$ or both are in $\widehat{V}^{(i+1)}$. Thus the result follows from [Lemma 8.21](#) by the triangle inequality.

We now sketch one of several ways of adapting this to the case when there are horoballs, leaving the details to the reader. We start with the $\widehat{B}$ geodesic joining x to y . Consider the points $x = x_0, x_1, x_2, \dots, x_k = y$ along this geodesic such that the $x_i, 0 < i < k$ are exactly the entry and exit points of the geodesic into the 1-neighborhoods of horoballs. As above, for $0 < i < k$, we can find $x''_i \in \widehat{B}^\theta$ very close to x_i , and such that $\pi_U(x''_i)$ is uniformly close to $\pi_U(x)$ for all angle domains U . (It is helpful to first note that a bounded length entry of a geodesic into a horoball cannot substantially change the angle projection associated to a hyperplane that that segment does not come close to. One can prove this by radially pushing to the boundary of the horoball and using that the angle projection is coarsely Lipschitz.)

It suffices to show that $d_{\widehat{B}^\theta}(x''_i, x''_{i+1})$ is bounded for all i .

In the first case, the geodesic from x_i to x_{i+1} stays in the thick part, and we argue exactly as in the cocompact case.

In the second case, the geodesic goes into (or at least close to) a horoball. First, we replace the geodesic with a path on the boundary of the 1-neighborhood of the horoball that is a geodesic in each factor $\widehat{\mathbb{C}}_i$. [Lemma 8.15](#) gives that the visual angle at each singular point in $\widehat{\mathbb{C}}_i$ bounded in terms of κ' , so factor by factor we can push slightly away from the singular points to get a thick path of bounded length. $\square$

Proof of [Proposition 11.1](#). Consider the geodesic $\gamma = [x, y]$ in $\widehat{B}$.

Case 1: The distance between x and y is large in $\widehat{B}$.

If all the regions of $\widehat{B}$ that are electrified when defining $\mathcal{C}(\widehat{B}^\circ)$ have small projection onto γ , the closest point projection from $\widehat{B}$ to γ extends to a coarsely Lipschitz map from the electrification of $\widehat{B}$ to the interval parameterizing γ , and thus the distance between x and y in $\widehat{B}$ is comparable to the distance in the electrification. Hence we can take $U = \widehat{B}^\circ$.

Otherwise, consider an electrified region that is minimal under inclusion which has large projection to γ . This implies that γ stays close to the electrified region for a long time, and hence the projection of x and y to the electrified region are far apart. If that region is a stratum, then repeating this argument in this stratum gives the result. If the region is a horoball the result follows from [Lemma 7.12](#).

Case 2: The distance between x and y is small in $\widehat{B}$. This case was handled in [Lemma 11.2](#). $\square$

11.2. Progress localization. We now address progress localization, building on the analysis showing enough projections and the analysis in [Section 10](#).

Proposition 11.3. *For any $P \geq 0$ there exists $N \geq 0$ and $R \geq 0$ such that for any V and any $x, y \in \widehat{B}^\theta$ with*

$$d_V(\pi_V(x), \pi_V(y)) \leq P$$

there is a set of at most N domains $U_i \subsetneq V$ such that if $T \subsetneq V$ is not nested in or orthogonal to one of the U_i then $d_T(\pi_T(x), \pi_T(y)) \leq R$.

Lemma 11.4. *There is a $D > 0$ such that the following holds. Let $\widehat{S}$ be a stratum closure, and let $x \in \widehat{B}^\theta$. There exists a set $\mathcal{E}$ of at most D center or angle domains such that the following holds.*

Consider the geodesic from x to the closest point projection $p_{\widehat{S}}(x)$ of x to $\widehat{S}$. If the geodesic $[x, p_{\widehat{S}}(x)]$ goes into a horoball $\widehat{T}_c$ that $\widehat{S}$ goes into, define x' to be the first entry into such a horoball. Otherwise set $x' = p_{\widehat{S}}(x)$. If U is nested in the stratum domain of $\widehat{S}^\circ$, and either $\pi_U(x')$ is undefined or

$$d_U(\pi_U(x), \pi_U(x')) \geq D,$$

then U is an angle domain, and either U is equal to one of the angle domains in $\mathcal{E}$ or U is orthogonal to one of the center domains in $\mathcal{E}$.

Recall from [Remark 4.3](#) that $\pi_U(x')$ is undefined when $x' \notin \widehat{B}^U$.

Proof. It is helpful to keep [Figure 9](#) in mind.

Case 1: If U is a stratum domain, then the corresponding stratum is contained in $\widehat{S}$, so [Corollary 10.3](#) gives that $d_U(\pi_U(x), \pi_U(x'))$ is uniformly small.

Case 2: If U is a cusp domain, then $\widehat{S}$ goes into the corresponding horoball, so [Corollary 10.3](#) gives that $d_U(\pi_U(x), \pi_U(x'))$ is uniformly small.

Case 3: If U is an angle domain corresponding to a hyperplane $\widehat{H}$, then $\widehat{H}$ must intersect $\widehat{S}$ but must not contain $\widehat{S}$. Assume $d_U(\pi_U(x), \pi_U(x'))$ is large, and apply [Lemma 10.4](#) to get a point x'' on $[x, x']$ that is either within r_θ of $\widehat{H}$ or in a horoball that $\widehat{H}$ enters but $\widehat{S}$ does not enter. By definition, x and x'' have the same projection to $\widehat{S}$.

We now claim that $\widehat{H} \cap \widehat{S}$ contains a point near $p_{\widehat{S}}(x)$. To see this, first suppose that x'' is within r_θ of $\widehat{H}$. In this case the claim follows from the fact that the projection of $\widehat{H}$ to $\widehat{S}$ is $\widehat{S} \cap \widehat{H}$ ([Remark 5.10](#)) and the fact that projection to $\widehat{S}$ is distance non-increasing, and the definition of “near” we get in this case can be taken to be “within r_θ ”. It remains to suppose that x'' is in a horoball that $\widehat{H}$ enters but $\widehat{S}$ does not. In this case the claim follows by additionally applying [Remark 4.2](#), and the definition of “near” we get in this case is “within 1”.

Case 3a: If $x' \neq p_{\widehat{S}}(x)$, then $p_{\widehat{S}}(x)$ is in some horoball $\widehat{T}_c$. Let $\mathcal{E}$ consist of only the center domain of $\widehat{T}_c$. Because $p_{\widehat{S}}(x)$ is at most 1 away from $\widehat{H} \cap \widehat{S}$, and we can assume (via $s_\theta \gg_{n, \mathcal{H}, \Gamma} 1$ in [Remark 4.2](#)) that any hyperplane that goes close to a horoball must enter it, we get that $\widehat{H}$ intersects $\widehat{S} \cap \widehat{T}_c$. This means that U is orthogonal to the center domain of the horoball $\widehat{T}_c$ containing $p_{\widehat{S}}(x)$.

Case 3b: We now assume that $x' = p_{\widehat{S}}(x)$, so $p_{\widehat{S}}(x)$ is not in a horoball. For this case, the following is helpful.

Remark 11.5. Suppose $\hat{H}$ is a hyperplane that enters the horoball of a cusp point $\hat{c}'$, and also that $\hat{H}$ intersects a stratum closure $\hat{S}$, and $\hat{S}$ does not enter the horoball of $\hat{c}'$. Then, by [Remark 5.10](#), $\hat{H}$ contains the projection $p_{\hat{S}}(\hat{c}')$. (The fact that $p_{\hat{S}}$ extends continuously to $\hat{c}'$ follows from a short comparison geometry argument and the fact that $p_{\hat{S}}$ is distance non-increasing.)

Fix $L > 0$ large enough so that the projection of any horoball onto another distance at least L away has diameter at most r_θ . Let y be the last point on $[x, x']$ that is not in a horoball and has distance at least L from x' ; or let $y = x$ if such a point does not exist.

Let $\mathcal{E}$ consist of all angle domains of hyperplanes that x' is within r_θ of but don't contain $\hat{S}$, and all angle domains associated to hyperplanes containing one of the points $p_{\hat{S}}(\hat{c}')$ but not containing $\hat{S}$, where the $\hat{c}'$ correspond to horoballs that $[y, x']$ enters.

The result now follows from [Lemma 10.4](#). Indeed, suppose $\hat{H}$ is a hyperplane that $[x, x']$ comes within r_θ of, or that enters a horoball that $[x, x']$ enters. If $[x, x']$ comes within r_θ of a hyperplane $\hat{H}$ that intersects $\hat{S}$, then $\hat{H} \cap \hat{S}$ is within r_θ of x' . Similarly if $[x, y]$ enters a horoball containing a hyperplane $\hat{H}$ that intersects $\hat{S}$, then $\hat{H} \cap \hat{S}$ must be extremely close to x' . Finally, if $[y, x']$ enters a horoball that $\hat{H}$ enters, then $\hat{H}$ contains one of the $p_{\hat{S}}(\hat{c}')$. $\square$

We also have a similar statement with the stratum replaced by a horoball.

Lemma 11.6. *There is a $D > 0$ such that the following holds. Let $\hat{T}_{\hat{c}}$ be a horoball, and let $x \in \hat{B}^\theta$. There exists a set $\mathcal{E}$ of at most D angle domains such that the following holds. Consider the geodesic from x to the closest point projection $x' = p_{\hat{T}_{\hat{c}}}(x)$ of x to $\hat{T}_{\hat{c}}$. If U is properly nested in a cusp domain associated to $\hat{T}_{\hat{c}}$ and either $\pi_U(x')$ is not defined or*

$$d_U(\pi_U(x), \pi_U(x')) \geq D,$$

then U is equal to one of the angle domains in $\mathcal{E}$.

Proof. The proof is similar to the proof of [Lemma 11.4](#) and is left to the reader. The only kind of domain that can be properly nested in a cusp domain is an angle domain, so one only needs to adapt Case 3. Since x' in this case is always $p_{\hat{T}_{\hat{c}}}(x)$, one does not have an analogue of Case 3a; one only has 3b. $\square$

Finally we can prove the progress localization assumption.

Proof of Proposition 11.3. There are only two types of V that are not nest-minimal, namely stratum domains and tree domains. Our goal is to build a set of domains $\{U_i\}$ with the desired properties.

Case 1: V is a stratum domain. Let $\hat{S}$ be the associated stratum closure. [Lemma 11.4](#) gives us points x', y' that are either in $\hat{S}$ or in a horoball that $\hat{S}$ enters. Either way, these points define points $x'', y'' \in \mathcal{C}(V)$; if $x' \in \hat{S}$ then we have $x'' = x'$ viewed as a point of $\mathcal{C}(V)$, and otherwise x'' is the generic point of the horoball in whose boundary it lies.

We now consider a geodesic (or path of approximately minimal length) α from x'' to y'' in $\mathcal{C}(V)$. By definition this can contain at most $d_V(x'', y'')/2 + 2$ generic points, and all the subpaths in $\hat{S}$ of course have length bounded by $d_V(x'', y'')$.

None of the subpaths can go more than depth 1 into a horoball, so there is no harm to assuming they are disjoint from horoballs, and we will do this. (The paths could be pushed out of horoballs.)

We construct the following set of domains $\{U_i\}$ consisting of the following.

- (1) The domains given by the use of [Lemma 11.4](#) above.
- (2) For each generic point of a horoball (intersected by $\hat{S}$) that α passes through, add the cusp domains associated to the horoball that are nested in V .
- (3) For each generic point of a codimension 1 hyperplane (in $\hat{S}$) that α contains, add the corresponding stratum domain (when the corresponding stratum has positive dimension), and the angle domain corresponding to the hyperplane containing that codimension 1 locus but not containing all of $\hat{S}$.
- (4) For each hyperplane not containing $\hat{S}$ that comes within r_θ of a bounded length subpath of α in $\hat{S}$, add the corresponding angle domain.

This set $\{U_i\}$ of domains is of bounded size (depending on P), so it suffices now to show that all domains $T \subsetneq V$ which are properly nested in V but not nested in or orthogonal to one of the $\{U_i\}$ have that $d_T(\pi_T(x), \pi_T(y))$ is bounded. Note that if $T \subsetneq V$ is not nested in or orthogonal to U_i , then $T \pitchfork U_i$ or $U_i \subsetneq T$.

By construction of $\{U_i\}$, we have that T cannot be a cusp domain associated to a horoball that α passes through, nor can it be an angle domain associated to a hyperplane that enters such a horoball.

Case 1a: T is a stratum domain or a cusp domain. In the stratum domain case, $\hat{B}^T = \hat{B}$, and in the cusp domain $\hat{B}^T$ is the complement of a horoball. We have that α consists of bounded subpaths in $\hat{S} \cap \hat{B}^T$ and visits to generic points. The bounded subpaths have uniformly bounded image in $\mathcal{C}(T)$ because π_T is coarsely Lipschitz on $\hat{B}^T$. The start and end points of an excursion to a generic point map to nearby points of $\mathcal{C}(T)$ by [Proposition 10.2](#).

Case 1b: T is an angle domain. Let $\hat{H}$ be the corresponding hyperplane. As mentioned above, $\hat{H}$ does not enter any horoball that α enters. Additionally, we have that α does not come within r_θ of $\hat{H}$. So α consists of subpaths in $\hat{S} \cap \hat{B}^T$ and generic points of regions contained in $\hat{B}^T$. The conclusion is the same as in Case 1a.

Case 2: V is a tree domain. Only angle domains can be nested in tree domains, and no angle domain can be nested in another.

Let x' and y' be the projections to the horoball associated to the tree domain. [Lemma 11.6](#) controls the passage from x to x' and from y' to y , so it suffices to think about how x' and y' differ. The result follows from [Proposition 8.14](#) and is a variant of [Corollary 8.23](#). $\square$

11.3. Proof of hierarchical hyperbolicity. Putting everything together, we get hierarchical hyperbolicity for our geometric model of the group.

Theorem 11.7. $(\hat{B}^\theta, \mathfrak{S}^b, \pi_\bullet)$ satisfies all the assumptions of [Proposition 3.6](#). In particular, $\hat{B}^\theta$ is an HHS.

Proof. In [Proposition 9.1](#) we checked that $(\hat{B}^\theta, \mathfrak{S}^b, \pi_\bullet)$ is a pre-HHS, in [Proposition 9.14](#) we checked that $\mathfrak{S}^b$ is a nearly valid index set, and in [Section 10](#) we

defined the active regions, so it suffices to check the numbered assumptions of [Proposition 3.6](#).

- (1) The enough projections assumption was checked in [Proposition 11.1](#).
- (2) The density assumption was checked in [Lemma 10.1](#).
- (3) The bounded projections assumption was checked in [Proposition 10.2](#) and [Lemma 10.5](#).
- (4) The structured redundancy assumption was checked in [Proposition 10.6](#) and [Proposition 10.7](#).
- (5) The progress localization assumption was checked in [Proposition 11.3](#).

Thus [Proposition 3.6](#) gives the result. $\square$

We can now conclude.

Proof of [Theorem 1.2](#). This now follows from [Proposition 3.7](#), since all of our constructions are natural and canonical enough to get the required equivariance under isometries. The condition on finitely many orbits is satisfied thanks to [Lemma 9.9](#) and the fact that there are only finitely many orbits of horoballs and strata. $\square$

11.4. The motivating application. For completeness, we now briefly reference the specific results that show that [Theorem 1.2](#) applies to give [Theorem 1.1](#).

Let $\mathcal{E}$ be the ring of Eisenstein integers, and let $\Lambda = \mathcal{E}^{4,1}$ be the free module $\mathcal{E}^5$ equipped with the hermitian form

$$h(x, y) = -x_0\overline{y_0} + x_1\overline{y_1} + \cdots + x_4\overline{y_4}.$$

Let $\Gamma \subset PU(4, 1)$ be the projectivized automorphism group of Λ , and let $\mathcal{H}$ be the union of the hyperplanes in B^4 defined using the perps of the norm 1 elements of Λ .

[\[ACT02a, Theorem 2.20\]](#) shows that $\mathcal{M}_{cs}$ is isomorphic to $\Gamma \backslash (B^4 - \mathcal{H})$. [\(A\)](#) is checked in [\[All00a, Proof of Corollary 3.4\]](#), [\(B\)](#) in [\[ACT02a, Lemma 7.28\]](#), and [\(C\)](#) is by definition. [\[ACT02a, Theorem 7.21, Theorem 8.2\]](#) shows that Γ has only one orbit of cusps.

One can verify [\(D\)](#) by noting that the null vector $(1, 1, 0, 0, 0)$ defines a cusp that is a limit of the 1-dimensional stratum defined by norm 1 elements

$$(0, 0, 1, 0, 0), (0, 0, 0, 1, 0), \text{ and } (0, 0, 0, 0, 1).$$

11.5. The Farrell-Jones Conjecture. What [\[DMS25\]](#) calls the Farrell-Jones Conjecture is what [\[BB19\]](#) calls the Farrell-Jones Conjecture with wreath products; it means that the wreath product with every finite group satisfies the K- and L-theoretic Farrell-Jones Conjectures with coefficients in any additive category. This version of the conjecture has good inheritance properties, many of which are summarized in [\[GR13, Proposition 1.1\]](#). What is important here is that the class of groups that satisfy the Farrell-Jones Conjecture is stable under taking subgroups, finite direct products, central extensions with finitely generated kernel, and passing to finite index overgroups, and that hyperbolic groups satisfy the Farrell-Jones Conjecture. We make use of these properties below without further comment.

We recall from [\[DMS25, Theorem C\]](#) that if $\widehat{\Gamma}$ is a colorable HHG, and if $\text{Stab}(U)$ satisfies the Farrell-Jones Conjecture for every non-nest maximal domain U , then $\widehat{\Gamma}$ satisfies the Farrell-Jones Conjecture. Colorability is a mild technical condition that is checked under the assumptions of [Theorem 1.2](#) in [Lemma B.15](#).

The core of the proof is the following two observations.

Lemma 11.8. *For any stratum closure $\hat{S} \subset \hat{B}$ of codimension k , $\text{Stab}_{\hat{\Gamma}}(\hat{S})$ has a finite index subgroup $\hat{\Gamma}_0$ that is a central $\mathbb{Z}^k$ -extension of a group as in [Theorem 1.2](#) for a ball of dimension $n - k$.*

In the degenerate case $k = n$ the stabilizer is virtually $\mathbb{Z}^n$.

Proof. Γ has a finite index torsion free subgroup Γ' , whose preimage $\hat{\Gamma}'$ is a finite index subgroup of $\hat{\Gamma}$. We set $\hat{\Gamma}_0$ to be the finite index subgroup of $\hat{\Gamma}' \cap \text{Stab}_{\hat{\Gamma}}(\hat{S})$ which does not permute the hyperplanes containing $\hat{S}$. This $\hat{\Gamma}_0$ acts on $\hat{S}$ with central kernel $\mathbb{Z}^k$, where k is the codimension of $\hat{S}$. The image of $\hat{\Gamma}_0$ in the isometry group of $\hat{S}$ is a group as in [Theorem 1.2](#) for a ball of dimension $n - k$. $\square$

Lemma 11.9. *For any horoball $\hat{T}_{\hat{e}}$, the stabilizer $\text{Stab}_{\hat{\Gamma}}(\hat{T}_{\hat{e}})$ has a finite index subgroup that is a product of $\mathbb{Z}$ and $n - 1$ finitely generated non-abelian free groups.*

Proof. Although it is not required, we prefer to use the discussion in [Remark 7.7](#). This remark shows that $\hat{T}_{\hat{e}}$ deformation retracts onto an $\mathbb{R}$ -bundle over a product of $n - 1$ finite valence trees. After passing to a finite index subgroup, one can assume that the image acts on the product of the $n - 1$ trees as a product of $n - 1$ cocompact free group actions. The kernel of this action is a central copy of $\mathbb{Z}$.

Because [Remark 7.7](#) gives that the monodromy of the $\mathbb{R}$ -bundle is trivial, one gets that the finite index subgroup is abstractly a subgroup of $\mathbb{R}$ times the $n - 1$ finitely generated free groups. When Q is a product of finitely generated free groups any central extension $1 \rightarrow \mathbb{Z} \rightarrow H \rightarrow Q \rightarrow 1$ that embeds in $\mathbb{R} \times Q$ in such a way that the maps to Q agree must be the trivial extension $\mathbb{Z} \times Q$. $\square$

Proof of [Theorem 1.3](#). We prove this by induction on the dimension n . For $n = 1$, it is true because $\hat{\Gamma}$ has a finite index subgroup which is a free group or a surface group. Assume the result is known for all dimensions less than some $n > 1$, and consider a $\hat{\Gamma}$ in dimension n .

We now wish to apply [\[DMS25, Theorem C\]](#). A minor subtlety is that $\mathfrak{S}^b$ is not the index set of the HHG structure on $\hat{\Gamma}$. The actual index set $\mathfrak{S}$ consists of tuples of orthogonal domains in $\mathfrak{S}^b$. This means that the stabilizer of a domain in $\mathfrak{S}$ is, up to finite index, an intersection of the stabilizers of a finite set of domains in $\mathfrak{S}^b$.

It follows that the stabilizers of domains in $\mathfrak{S}$, up to finite index, are all of the form $\text{Stab}_{\hat{\Gamma}}(\hat{S})$ or $\text{Stab}_{\hat{\Gamma}}(\hat{T}_{\hat{e}})$ or $\text{Stab}_{\hat{\Gamma}}(\hat{S}) \cap \text{Stab}_{\hat{\Gamma}}(\hat{T}_{\hat{e}})$, where in this final case the stratum closure $\hat{S}$ must enter the horoball $\hat{T}_{\hat{e}}$.

We see that $\text{Stab}_{\hat{\Gamma}}(\hat{S})$ satisfies the Farrell-Jones Conjecture by [Lemma 11.8](#) and the induction hypothesis, and $\text{Stab}_{\hat{\Gamma}}(\hat{T}_{\hat{e}})$ satisfies the Farrell-Jones Conjecture by [Lemma 11.9](#). That $\text{Stab}_{\hat{\Gamma}}(\hat{S}) \cap \text{Stab}_{\hat{\Gamma}}(\hat{T}_{\hat{e}})$ satisfies the Farrell-Jones Conjecture is then automatic because it is a subgroup of a group that satisfies the Farrell-Jones Conjecture. The result now follows from [\[DMS25, Theorem C\]](#). $\square$

Remark 11.10. A related analysis shows that the HHG structures provided by [Theorem 1.2](#) have cobounded product regions in the sense of [\[Man25, Definition 1.13\]](#).

APPENDIX A. PROOF OF CRITERIA FOR HIERARCHICAL HYPERBOLICITY

A.1. The definition of an HHS. Here we follow the definition in [BHS19, Section 1.3]. We have made minor changes to the definition, which do not change which spaces the definition applies to. These changes are noted after the definition. We continue to use “asynchronous” as a replacement for the standard term “transverse”.

Definition A.1. Let $(\mathcal{X}, \mathfrak{S}, \pi_\bullet)$ be a pre-HHS with $\mathfrak{S}$ a valid index set. We say that $(\mathcal{X}, \mathfrak{S}, \pi_\bullet)$ is a hierarchically hyperbolic space if the following hold for some constants $\kappa_0, \alpha \geq 0, E \geq \kappa_0$.

- (1) **Nesting:** For each properly nested pair of domains $U \subsetneq V$ there is a specified point $\rho_V^U \in \mathcal{C}(V)$.
- (2) **Transversality:**
 - (a) If $V \triangleleft W$, then there are specified points $\rho_W^V \in \mathcal{C}(W)$ and $\rho_V^W \in \mathcal{C}(V)$ such that for all $x \in \mathcal{X}$

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

- (b) If $U \subsetneq V$ and both ρ_W^U and ρ_W^V are defined, then $d_W(\rho_W^U, \rho_W^V) \leq \kappa_0$.

- (3) **Large links:** For any $P \geq 0$ there exists $N \geq 0$ such that for any V and any $x, y \in \mathcal{X}$ with

$$d_V(\pi_V(x), \pi_V(y)) \leq P$$

there is a set of at most N domains $U_i \subsetneq V$ such that if $T \subsetneq V$ is not nested in one of the U_i then $d_T(\pi_T(x), \pi_T(y)) \leq E$.

- (4) **Bounded geodesic image:** For all $x, y \in \mathcal{X}$ and $U \subsetneq V$, if there is a geodesic in $\mathcal{C}(V)$ from $\pi_V(x)$ to $\pi_V(y)$ that stays $E + 2\delta$ away from ρ_V^U , then

$$d_U(\pi_U(x), \pi_U(y)) \leq E.$$

Here δ is such that all $\mathcal{C}(U)$ are δ -hyperbolic.

- (5) **Partial realization:** For any orthogonal family $\{V_j\}$ and points $p_j \in \mathcal{C}(V_j)$, there exists a point $x \in \mathcal{X}$ such that

$$d_{V_j}(\pi_{V_j}(x), p_j) \leq \alpha$$

for all j , and such that for each j and each V with $\rho_V^{V_j}$ defined,

$$d_V(\pi_V(x), \rho_V^{V_j}) \leq \alpha.$$

- (6) **Uniqueness:** For all $\kappa \geq 0$ there exists $\omega \geq 0$ such that if $d(x, y) \geq \omega$ then

$$d_U(\pi_U(x), \pi_U(y)) \geq \kappa$$

for some $U \in \mathfrak{S}$.

The minor changes we have made are the following: we have chosen to use points everywhere instead of sets of uniformly bounded diameter following the discussion in [Sis19, Section 3]; we have incorporated the coarse surjectivity assumption of [BHS19, Section 1.3] into the partial realization axiom; and we have modified the large links axiom following [BHS19, Remark 1.5] and [Rus22, Remark 2.10].

Note that “ ρ_V^U is defined” means that either $U \subsetneq V$ or $U \triangleleft V$.

A.2. The large links axiom. Consider the following less uniform version of the large links axiom, where the uniform constant E has been replaced with a constant R that can depend on P .

- (3') **Less uniform large links:** For any $P \geq 0$ there exists $N \geq 0$ and $R \geq 0$ such that for any V and any $x, y \in \mathcal{X}$ with

$$d_V(\pi_V(x), \pi_V(y)) \leq P$$

there is a set of at most N domains $U_i \subsetneq V$ such that if $T \subsetneq V$ is not nested in one of the U_i then $d_T(\pi_T(x), \pi_T(y)) \leq R$.

We remark that this change in uniformity is not significant.

Lemma A.2. *In the presence of the other axioms, the less uniform large links axiom implies the usual large links axiom, possibly after increasing the constant E .*

The proof does not use hyperbolicity of $\mathcal{C}(U)$ when U is nest minimal.

Proof. Fix even constants $P_0 \gg L \gg \max(E, \delta, \alpha)$. Apply the less uniform large links with P_0 to obtain an N_0 and an R_0 as above. We will now show that the usual large links axiom holds if E is increased to $E' = 2E + R_0$.

Let $P > 0$ be arbitrary, and suppose $d_V(\pi_V(x), \pi_V(y)) \leq P$. If $d_V(\pi_V(x), \pi_V(y)) \leq P_0$ there is nothing new to prove, so assume $d_V(\pi_V(x), \pi_V(y)) > P_0$.

Fix a geodesic from $\pi_V(x)$ to $\pi_V(y)$. Pick points $p_0, \dots, p_d$ on this geodesic, so $d_V(\pi_V(x), p_i) = i$ and $d_V(p_d, \pi_V(y)) < 1$. Using partial realization, pick points

$$x = x_0, x_1, \dots, x_d = y$$

in $\mathcal{X}$ so that $\pi_V(x_i)$ lies within α of p_i .

For each valid pair $i, i + L$ of indices, apply the less uniform large links to the points x_i, x_{i+L} to get a set $\mathcal{U}_i$ of at most N_0 domains properly nested in V . We claim that we can use the set $\mathcal{U} = \bigcup_i \mathcal{U}_i$, which has size at most dN_0 , to prove the usual version of large links. We use $N = (P + 1)N_0$, which is greater than dN_0 and does not depend on the choice of x, y .

Indeed, suppose $T \subsetneq V$ and T is not nested in any of the domains of $\mathcal{U}$. If the chosen geodesic from $\pi_V(x)$ to $\pi_V(y)$ does not come within $E + 2\delta$ of ρ_V^T , the bounded geodesic image axiom shows that $d_T(\pi_T(x), \pi_T(y)) \leq E < E'$. So suppose that some p_j is within $E + 2\delta + 1$ of ρ_V^T .

We leave the case when p_j is close to the beginning or end of the geodesic to the reader. Otherwise, let $i = j - L/2$, and consider x_i and x_{i+L} . If $k \leq i$ or $k \geq i + L$, then the triangle inequality gives $d_V(\pi_V(x_k), \rho_V^T) \gg E + 2\delta$. Hence the bounded geodesic image axiom gives that $\pi_T(x_i)$ is within E of $\pi_T(x)$ and $\pi_T(x_{i+L})$ is within E of $\pi_T(y)$. This uses hyperbolicity of $\mathcal{C}(V)$, to ensure that a geodesic from $\pi_V(x_i)$ to $\pi_V(x)$, and a geodesic from $\pi_V(x_{i+L})$ to $\pi_V(y)$, stays close to the appropriate segments of the chosen geodesic from $\pi_V(x)$ to $\pi_V(y)$. Since T is not nested in any domain in $\mathcal{U}_i$, we get that $d_T(\pi_T(x_i), \pi_T(x_{i+L})) \leq R_0$. So the triangle inequality gives

$$d_T(\pi_T(x), \pi_T(y)) \leq E + R_0 + E = E'$$

as desired. $\square$

A.3. The proof of Proposition 3.6. We start by explaining how to add domains to a nearly valid index set to make it into a valid index set. This can be compared to, but is easier and less general than, the appendix of [ABD21]. (We cannot use the appendix of [ABD21] because of [ABR25, Remark 3.4].)

Definition A.3. If $\mathfrak{S}^b$ is a nearly valid index set, we define $\mathfrak{S}$ to be the set of non-empty orthogonal subsets of $\mathfrak{S}^b$. We say $\mathbf{U}, \mathbf{V} \in \mathfrak{S}$ are orthogonal if every domain in $\mathbf{U}$ is orthogonal to every domain in $\mathbf{V}$, and we say $\mathbf{U}$ is nested in $\mathbf{V}$ if every domain in $\mathbf{U}$ is nested in some domain of $\mathbf{V}$. As always, two domains in $\mathfrak{S}$ are asynchronous if they are not orthogonal and neither is nested in the other.

Example A.4. Say $U_1 \sqsubset V$ and $V \perp U_2$ are domains in $\mathfrak{S}^b$. Then $\mathbf{U} = \{U_1, U_2\}$ is asynchronous to $\mathbf{V} = \{V\}$ (even though no domain in $\mathbf{U}$ is asynchronous to any domain in $\mathbf{V}$).

We view $\mathfrak{S}^b$ as a subset of $\mathfrak{S}$ via the inclusion $U \mapsto \{U\}$, and note that the orthogonality and nesting relations on $\mathfrak{S}$ extend those on $\mathfrak{S}^b$. We will typically use regular letters like U for elements of $\mathfrak{S}^b$ and bold letters like $\mathbf{U}$ for elements of $\mathfrak{S}$, except that sometimes, keeping the inclusion in mind, we may write U for $\{U\}$.

The point of this construction is the following.

Lemma A.5. *If $\mathfrak{S}^b$ is a nearly valid index set then $\mathfrak{S}$ is a valid index set.*

Proof. It is immediate that the orthogonality relation on $\mathfrak{S}$ is symmetric and anti-reflexive, and that the nesting relation is a partial order. If $\Sigma \in \mathfrak{S}^b$ is nest maximal, then $\{\Sigma\} \in \mathfrak{S}$ is nest maximal.

The finite height axiom follows from the corresponding statement for $\mathfrak{S}^b$ and the fact that orthogonal sets of domains in $\mathfrak{S}^b$ have bounded size. The coherence axiom is left as an exercise, so it only remains to check the containers axiom.

Suppose $\mathbf{U} = \{U_i\}$ is nested in $\mathbf{V} = \{V_j\}$. First observe that it suffices to construct containers when $\mathbf{V} = \{V\}$ has size 1, since containers for $\mathbf{U}$ in the various V_j can be assembled into a container for $\mathbf{U}$ in $\mathbf{V}$.

To build a container for $\{U_1, \dots, U_r\}$ in $\{V\}$, start by applying the weak container axiom to the pair U_1, V , to find an orthogonal tuple of domains $\{C_1, \dots, C_s\}$ properly nested in V . Note that each $U_i, i > 1$ is nested in a unique C_j , so one can iterate the procedure. A stronger fact is proven in [Lemma B.5](#) so here we leave the details to the reader. $\square$

With this in hand, we can proceed to the proof.

Proof of Proposition 3.6. We let $\mathfrak{S}$ be as in [Definition A.3](#), so $\mathfrak{S}$ is a valid index set by [Lemma A.5](#). For $\mathbf{U} \in \mathfrak{S}$, we define $\mathcal{C}(\mathbf{U})$ and $\pi_{\mathbf{U}}$ as follows: If $\mathbf{U} = \{U\}$, then $\mathcal{C}(\mathbf{U}) = \mathcal{C}(U)$ and $\pi_{\mathbf{U}} = \pi_U$, and otherwise define $\mathcal{C}(\mathbf{U})$ to be a point and $\pi_{\mathbf{U}}$ to be the constant map. It now suffices to check the 6 axioms of [Definition A.1](#). We will view $\mathfrak{S}^b$ as a subset of $\mathfrak{S}$ in what follows.

We will need points $\rho_{\mathbf{V}}^{\mathbf{U}} \in \mathcal{C}(\mathbf{V})$ when $\mathbf{U} \triangleleft \mathbf{V}$ or $\mathbf{U} \sqsubset \mathbf{V}$. If $\mathbf{V} \notin \mathfrak{S}^b$, then $\mathcal{C}(\mathbf{V})$ is a point, and $\rho_{\mathbf{V}}^{\mathbf{U}}$ is that point. So assume $\mathbf{V} = \{V\}$ is a singleton, and $\mathbf{U} = \{U_i\}$. To have definitions that are coarsely natural, we first note the following.

Lemma A.6. *Suppose $U_1 \perp U_2$, and each U_i is properly nested in or asynchronous to V . Then $\text{diam } \pi_V(\mathcal{A}_{U_1} \cup \mathcal{A}_{U_2}) \leq 2B$.*

Proof of Lemma A.6. The density assumption gives $\mathcal{A}_{U_1} \cap \mathcal{A}_{U_2} \neq \emptyset$ so the result follows from the bounded projections assumption. $\square$

First suppose $\mathbf{U} = \{U_i\}$ is nested in $\mathbf{V} = \{V\}$. Define $\rho_{\mathbf{V}}^{\mathbf{U}}$ to be any point in $\pi_V(\cup \mathcal{A}_{U_i})$. Since this set has diameter at most $2B$ by [Lemma A.6](#), this is coarsely well defined.

Next, suppose $\mathbf{U} = \{U_i\}$ is asynchronous to $\mathbf{V} = \{V\}$. We first claim that at least one U_i must be properly nested in or asynchronous to V . Otherwise, for each i either $V \subseteq U_i$ or $V \perp U_i$, and this would give the contradiction that either $\mathbf{V} \subseteq \mathbf{U}$ (if $V \subseteq U_i$ for some i) or $\mathbf{V} \perp \mathbf{U}$. We let S be the set of i with U_i properly nested in or asynchronous to V and define $\rho_{\mathbf{V}}^{\mathbf{U}}$ to be any point in $\pi_V(\cup_{i \in S} \mathcal{A}_{U_i})$. Since this set has diameter at most $2B$ by Lemma A.6, this is coarsely well defined. We can now proceed.

Proof of the transversality axiom. For part (a), we first note that the inequality is trivial unless $\mathbf{U} = \{U\}$ and $\mathbf{V} = \{V\}$ are both singletons, so assume this is the case. The result is now given by part of the structured redundancy assumption.

Part (b) follows from the final part of the bounded projections assumption. $\square$

Proof of the bounded geodesic image axiom. For this to fail, we must have that $d_{\mathbf{U}}(\pi_{\mathbf{U}}(x), \pi_{\mathbf{U}}(y))$ is large and $\rho_{\mathbf{V}}^{\mathbf{U}}$ is far from a geodesic from $\pi_{\mathbf{V}}(x)$ to $\pi_{\mathbf{V}}(y)$. So it suffices to consider the case when both domains $\mathbf{U} = \{U\}$, $\mathbf{V} = \{V\}$ are in $\mathfrak{S}^b$. The result is now given by part of the structured redundancy assumption. $\square$

Proof of the partial realization axiom. This follows immediately from the density assumption. $\square$

Proof of the large links axiom. We prove the weaker version, which is sufficient by Lemma A.2.

Fix $P \geq 0$. The progress localization assumption gives that there is a N_{PL} and an $R > 0$ such that for any $V \in \mathfrak{S}^b$ and any $x, y \in \mathcal{X}$ with $d_V(\pi_V(x), \pi_V(y)) \leq P$ there is a set of at most N_{PL} domains $U_i \subsetneq V$ such that if $T \subsetneq V$ is not nested in or orthogonal to one of the U_i then $d_T(\pi_T(x), \pi_T(y)) \leq R$.

We claim that the statement in large links holds with $N = \max(2N_{PL}, K)$, where K is the maximum size of a set of pairwise orthogonal domains. Indeed, consider $\mathbf{V}$ and $x, y \in \mathcal{X}$ with $d_{\mathbf{V}}(\pi_{\mathbf{V}}(x), \pi_{\mathbf{V}}(y)) \leq P$.

First consider the case when $\mathbf{V} = \{V\}$ is a singleton. Consider the set of domains $U_i \subsetneq V$ given by progress localization. Adding the container domains for the U_i in V , as described in Index Set Axiom 3, gives the required set of at most N domains.

Next consider the case when $\mathbf{V} = \{V_1, \dots, V_k\}$ is not a singleton. Then we simply define $\{U_i\}$ to be the set of the V_i , now viewed as domains in $\mathfrak{S}$ via the usual implicit map $V_i \mapsto \{V_i\}$. $\square$

This concludes the proof, because the uniqueness axiom is the same as the enough projections assumption. $\square$

A.4. The definition of an HHG. Suppose we have a self-map $f : \mathcal{X} \rightarrow \mathcal{X}$ of an HHS, where we use the notation above for the HHS. Suppose K is a constant. Suppose we have a bijection

$$f^\diamond : \mathfrak{S} \rightarrow \mathfrak{S}$$

preserving nesting, orthogonality, and transversality, and, for each $U \in \mathfrak{S}$ we have an isometry

$$f^*(U) : \mathcal{C}(U) \rightarrow \mathcal{C}(f^\diamond(U)).$$

Suppose that, for all $x \in \mathcal{X}$,

$$d_{f^\diamond(U)}(\pi_{f^\diamond(U)}(f(x)), f^*(U)(\pi_U(x))) \leq K,$$

and, whenever U is nested in or asynchronous to V , then

$$d_{f^\diamond(V)}(f^*(V)\rho_V^U, \rho_{f^\diamond(V)}^{f^\diamond(U)}) \leq K.$$

Then we say that f , together with data of the associated maps, is an automorphism of the HHS.

There is a natural notion of composing automorphisms, which involves composing the associated data as well as the actual self-maps f . We say that a group acts on an HHS if there is a homomorphism from the group to the group of automorphisms of the HHS.

Following [BHS19, Section 1.7], to prove that a group is hierarchically hyperbolic, it suffices to prove that it acts properly and cocompactly by isometries on an HHS and that the induced action on $\mathfrak{S}$ has finitely many orbits. (The actual definition allows for quasi-actions in addition to actions, and allows for metrically proper and cobounded actions in addition to actions that are proper and cocompact by isometries.)

We note that [BHS19, Section 1.7] does not use the variant axioms of [BHS19, Section 1.3] that we follow, but the proof of the equivalence in [BHS19, Proposition 1.11] can easily be adapted to bridge this difference.

A.5. Proof of Proposition 3.7. The proof is almost immediate. The group G acts on the enlarged index set $\mathfrak{S}$ via

$$g^\diamond\{U_1, \dots, U_k\} = \{g^\diamond U_1, \dots, g^\diamond U_k\}.$$

For a domain V in $\mathfrak{S} - \mathfrak{S}^b$, for any $g \in G$ we have that both $\mathcal{C}(V)$ and $\mathcal{C}(g^\diamond V)$ are points, and the required isometry $g(V) : \mathcal{C}(V) \rightarrow \mathcal{C}(g^\diamond V)$ is the constant map.

The requirement that

$$d_{f^\diamond(V)}(f^*(V)\rho_V^U, \rho_{f^\diamond(V)}^{f^\diamond(U)}) \leq K$$

follows from the definition of the ρ points and Lemma A.6.

Remark A.7. Although the criteria given are only sufficient conditions for hierarchical hyperbolicity, it seems straightforward to modify them to actually characterize hierarchical hyperbolicity. To show an appropriately modified criterion can apply to an arbitrary HHS, one might define the active regions to be the product regions of [BHS19], and apply [RST23, Proposition 4.24], which is a corrected form of a result in [BHS19]. If $U \subsetneq V$, to obtain a coarser form of $\mathcal{A}_U \cap \mathcal{A}_V \neq \emptyset$, one might want to note that the product region of U contains a point not far from a point of the product region of V , which (following a suggestion of Sisto) might be done by showing that for any point of the HHS the result of projecting (“gating”) it to the product region of U and then to the product region of V is close to the result of applying the projections in the opposite order.

APPENDIX B. THE COMBINATORICS OF THE INDEX SET (NOT REQUIRED)

In this appendix, which is not used in the paper, we explore the combinatorial properties of the sets $\mathfrak{S}^b$ and $\mathfrak{S}$ together with the relations of nesting and orthogonality. For example, as noted in Remark B.6, the set $\mathfrak{S}$ with the nesting partial order is, after adding a smallest element, a lattice (a poset with greatest lower bounds and least upper bounds). This appendix shows that the HHG structure we produce is extremely similar to the usual HHG structure on mapping class groups, and allows us to upgrade Theorem 1.2 as follows.

Theorem B.1. *The HHG structure produced in the proof of Theorem 1.2 is colorable and has wedges, clean containers, and the orthogonals for non-split domains property.*

The definitions of these terms will be recalled below.

Proof. Colorability is checked in Lemma B.15; wedges in Lemma B.4; clean containers in Lemma B.5; and the orthogonals for non-split domains property in Lemma B.10. $\square$

Remark B.2. It follows from Theorem B.1 and [HMS26, Theorem 1] that our spaces are Combinatorial Hierarchically Hyperbolic Spaces (CHHS) as defined in [BHMS24]. We believe it would require new ideas to check this directly without reproducing much of our analysis as well as parts of the proof of [HMS26, Theorem 1.1]. Additionally, an approach based on [BHMS24] would not allow our extension to the relative setting in Appendix C (see [HMS26, Remark 3.17]), and would not provide as much explicit information on the HHS structure.

Recall that $\mathfrak{S}^b$ is the set of stratum domains of positive dimensional strata, angle domains, center domains, and tree domains; $\mathfrak{S}$ is the set of subsets $\{U_1, \dots, U_k\}$ of $\mathfrak{S}^b$ consisting of $k \geq 1$ pairwise orthogonal domains of $\mathfrak{S}^b$; and we can view $\mathfrak{S}^b$ as a subset of $\mathfrak{S}$ via the inclusion $U \mapsto \{U\}$.

This appendix assumes familiarity with previously introduced material, especially Figure 8 and Lemma A.5.

B.1. Wedges. Given a poset, the wedge of two elements is their greatest lower bound, if it exists. We will say that a poset has wedges if every pair of elements either has a wedge or there does not exist any element less than both of them. (This is the same as saying wedges always exist after adding a smallest element.) We will show that $\mathfrak{S}$ has wedges, but first we need the following.

Lemma B.3. *Given $U, V \in \mathfrak{S}^b$, there is a unique subset $\{W_1, \dots, W_k\} \subset \mathfrak{S}^b$ of pairwise orthogonal domains such that*

- (1) *all W_i are nested in both U and V , and*
- (2) *every domain nested in both U and V is nested in one of the W_i .*

We call $\{W_1, \dots, W_k\}$ a wedge set for U and V . We allow $k = 0$, and note that the wedge set is empty exactly when there are no domains nested in both U and V .

Proof. It suffices to consider pairs U, V with neither nested in the other, since if $U \subseteq V$ then U is the wedge of U and V . We also assume that there is something nested in both U and V . We proceed in cases according to the type of U and the type of V . There are only three cases.

Case 1: stratum, stratum. The domains nested in a stratum domain are: smaller stratum domains; angle domains of orthogonally intersecting hyperplanes; center domains of horoballs the stratum enters, and some tree domains of horoballs the stratum enters (those not equal to a point in the tree domain factor).

Subcase 1a: The closures of the two strata intersect. In this case, let W be the stratum domain associated to the generic stratum of the intersection of the closures; the analysis below, using the assumption that there is something nested in both

U and V , will show that this intersection is not a point. We claim that W is the wedge of U and V , or in other words that $\{W\}$ satisfies the requirements of the lemma.

Now suppose that R is a domain nested in U and V . We want to show that R is nested in W . Considering U, V as strata, we can consider their closures $\overline{U}, \overline{V}$ and proceed as follows.

- (1) If R is a stratum domain, then $\overline{R} \subset \overline{U}$ and $\overline{R} \subset \overline{V}$ and it suffices to note that $\overline{R} \subset \overline{U} \cap \overline{V}$.
- (2) If R is an angle domain, then the associated hyperplane intersects both $\overline{U}$ and $\overline{V}$ orthogonally. We know $\overline{U}$ and $\overline{V}$ intersect orthogonally along $\overline{W}$. [Lemma 2.6](#) and [Remark 5.10](#) gives that the hyperplane associated to R intersects $\overline{W}$.
- (3) If R is a center domain, then both U and V enter the associated horoball. [Lemma 2.6](#) gives that W also enters the horoball.
- (4) If R is a tree domain, then both U and V enter the associated horoball, and again [Lemma 2.6](#) gives that W also enters the horoball.

In all cases one can deduce that R is nested in W .

Subcase 1b: Both strata enter a horoball, but their closures $\overline{U}$ and $\overline{V}$ do not intersect. First note that there cannot be more than one horoball that both strata enter, since otherwise the geodesic between the associated cusps would show that $\overline{U} \cap \overline{V} \neq \emptyset$, giving a contradiction.

Here are the domains nested in both U and V :

- (1) the center domain of the cusp,
- (2) each tree domain of this cusp associated to a factor $\hat{\mathbb{C}}_i$ where both U and V are not constant, and
- (3) angle domains nested in those tree domains.

The wedge set of U and V is the center and tree domains above.

Case 2: stratum, tree. Only angle domains can be properly nested in tree domains, so there is an angle domain $N^1(\hat{H})$ nested in both U and V . The $\hat{H}$ intersects $\overline{U}$ orthogonally, and corresponds to a point in the tree corresponding to the tree domain V .

Subcase 2a: The stratum enters the horoball associated to V . In this case the stratum is given by fixing the point coordinate in some of the tree factors – but, since U is not a subset of $\hat{H}$, not the factor corresponding to the given tree domain. In this case the given tree domain is actually nested in the given stratum domain. So this case does not occur because of our assumption that neither of U, V is nested in the other.

Subcase 2b: The stratum U does not enter the horoball associated to V . We know $N^1(\hat{H})$ is nested in U and V . Suppose another angle domain $N^1(\hat{H}')$ is nested in U and V . We claim that $\hat{H}$ and $\hat{H}'$ cannot be disjoint. Otherwise we can contradict the fact that closest point projection is distance non-increasing, because $\hat{H} \cap \overline{U}$ is the closest point projection of $\hat{H}$ to $\overline{U}$ ([Remark 5.10](#)) and similarly for $\hat{H}'$, and because $\hat{H}$ and $\hat{H}'$ enter the same horoball and hence get arbitrarily close.

Thus the collection of domains nested in U and V is a set of pairwise orthogonal angle domains, and this set is hence the wedge set.

Case 3: tree, tree. To have something nested in both it has to be an angle domain. Since we have assumed U and V are distinct, this means that the two tree domains are associated with different cusps. There is a unique complex line joining any two cusps, and any hyperplane entering both horoballs must contain this line. There is a unique hyperplane that both contains this line and is contained in the correct parallelism class at each cusp, and the wedge set should be the associated single angle domain. $\square$

Lemma B.4. $\mathfrak{S}$ has wedges.

Proof. Using Lemma B.3, we claim the wedge of $\{V_i\}$ and $\{U_j\}$ is the union $\{W_k\}$ over i, j of the wedge set of V_i and U_j .

First, we show that $\{W_k\} \in \mathfrak{S}$. Suppose W_{k_1}, W_{k_2} are in this $\{W_k\}$ and are distinct. We must show they are orthogonal. Say W_{k_1} is in the wedge set of V_{i_1} and U_{j_1} and W_{k_2} is in the wedge set of V_{i_2} and U_{j_2} . If $V_{i_1} \neq V_{i_2}$ or $U_{j_1} \neq U_{j_2}$ then the coherence condition in Definition 3.3 gives that $W_{k_1} \perp W_{k_2}$. In the remaining case we get $W_{k_1} \perp W_{k_2}$ from the definition of a wedge set.

Next, we show that $\{W_k\}$ is nested in both $\{V_i\}$ and $\{U_j\}$. Fix k , and say W_k is part of the wedge tuple for V_i and U_j . So W_k is nested in V_i and U_j . Since this is true for all k this gives the desired result.

Finally, we consider a $\{Q_r\} \in \mathfrak{S}$ that is nested in $\{V_i\}$ and $\{U_j\}$ and show $\{Q_r\}$ is nested in $\{W_k\}$. Fix r . We know Q_r is nested in some V_i and some U_j . So Q_r is nested in an element of the wedge tuple of V_i and U_j . Since this is true for all r this gives the desired result. $\square$

B.2. Clean containers. We now show the following.

Lemma B.5. Suppose $\{U_i\}, \{V_j\} \in \mathfrak{S}$ and $\{U_i\}$ is nested in $\{V_j\}$. Suppose there is at least one domain nested in $\{V_j\}$ and orthogonal to $\{U_i\}$. Then there is a domain $\{W_k\} \in \mathfrak{S}$ that is nested in $\{V_j\}$ and orthogonal to $\{U_i\}$, and such that every other such domain is nested in $\{W_k\}$.

One thinks of $\{W_k\}$ as the orthogonal complement of $\{U_i\}$ in $\{V_j\}$. The statement of Lemma B.5 is what is known in the HHG literature as $\mathfrak{S}$ having clean containers [ABD21, Definition 7.1].

Proof. First we prove this when $\{V_j\} = \{V\}$ is a singleton. The clean container in this case is the wedge (over i) of the orthogonal tuples produced by Lemma 9.13 for each $U_i \subsetneq V$.

With this having been proven, it suffices to note in general that the desired clean container is the orthogonal set obtained as the union over j_0 of the clean container for $\{U_i\}$ in the singleton $\{V_{j_0}\}$. $\square$

Remark B.6. As pointed out in [HMS26, Remark 10.21], the existence of wedges together with the finite height assumption in Definition 3.3 implies the existence of least upper bounds (joins). Moreover, $\mathfrak{S}$ with one smallest element added is a complete lattice.

B.3. Orthogonals for non-split domains. We start in the context of $\mathfrak{S}^b$. Recall that angle and center domains are always nest-minimal. The remaining domains have attached to them a certain very special set of domains that we call a halo.

- (1) If U is a stratum domain, define its halo ∂U to be the set of angle domains corresponding to hyperplanes containing U . If U has codimension k , this set has size k .
- (2) If U is a tree domain, define its halo ∂U to be the set consisting only of the associated center domain.

In either case, we note that ∂U consists of nest-minimal domains that are pairwise orthogonal. Except if U is the nest maximal domain, ∂U is non-empty when it is defined. We do not define the halo for angle and center domains.

Halos have the following two remarkable properties.

Lemma B.7. *Suppose U is a tree domain or a stratum domain, and suppose $U \subsetneq V$. Then at least one element of ∂U is nested in V .*

Proof. If U is a tree domain, then V is a (particular kind of) stratum domain entering the horoball of U , so the associated center domain is nested in V .

If U is a stratum domain, then V is a bigger stratum domain. There exists a hyperplane that contains U but not V , and the associated angle domain is nested in V . $\square$

Lemma B.8. *Suppose U is a tree domain or a stratum domain, and suppose that $H \in \partial U$ and $V \perp U$. Then $V = H$ or $V \perp H$.*

Proof. If U is a tree domain, then V is an angle, center, or tree domain, and H is the center domain. If U is a stratum domain, then V is an angle domain and H is an angle domain. Either way, the result follows from examining the orthogonality relation. $\square$

We now turn to $\mathfrak{S}$. Recall from [HMS26, Definition 3.6] that a $\{U_i\} \in \mathfrak{S}$ is split if there exists $\{V_i\}$ nested in $\{U_i\}$ such that every domain nested in $\{U_i\}$ is either orthogonal to $\{V_i\}$ or has $\{V_i\}$ nested in it.

Remark B.9. If one of the U_i , say U_1 , is nest-minimal, then $\{U_i\}$ is split, as seen by taking $\{V_i\} = \{U_1\}$.

Recall from [HMS26, Definition 3.9] that $\mathfrak{S}$ has the orthogonals for non-split domains property if whenever $\{U_i\} \subsetneq \{W_k\}$ and $\{U_i\}$ is not split, then there exists a domain nested in $\{W_k\}$ and orthogonal to $\{U_i\}$.

Lemma B.10. *$\mathfrak{S}$ has the orthogonals for non-split domains property.*

Proof. Suppose $\{U_i\} \subsetneq \{W_k\}$ and $\{U_i\}$ is not split. If there is some k_0 such that W_{k_0} does not have any of the U_i nested in it, it suffices to note that the singleton $\{W_{k_0}\}$ is orthogonal to $\{U_i\}$. Otherwise, without loss of generality $U_1 \subsetneq W_1$. By Remark B.9, U_1 is not nest-minimal. Lemma B.7 gives a $Q \in \partial U_1$ that is nested in W_1 . This Q is nest-minimal so it cannot be equal to any of the U_i , and it suffices to note that the singleton $\{Q\}$ is orthogonal to $\{U_i\}$. $\square$

B.4. Colorability. Recall that an HHG is called colorable if each domain can be colored by one of finitely many colors in such a way that domains of the same color are asynchronous and such that the action of the HHG on the set of domains induces an action of the HHG on the set of colors [Hag23].

Lemma B.11. *To prove $\mathfrak{S}$ is colorable it suffices to prove that $\mathfrak{S}^b$ is colorable.*

Proof. Suppose $\mathfrak{S}^b$ is colorable by a set of colors $\mathcal{C}$. We then color $\mathfrak{S}$ with a new set of colors equal to the power set $2^{\mathcal{C}}$. The color of $\{U_i\}$ is the set of colors of the U_i .

Suppose $\{U_i\} \neq \{V_j\}$, and assume without loss of generality that V_1 is not equal to any of the U_i . If $\{U_i\}$ has the same color as $\{V_j\}$, then without loss of generality U_1 has the same color as V_1 , which means that U_1 and V_1 are asynchronous. It follows that $\{U_i\}$ and $\{V_j\}$ are asynchronous. $\square$

Lemma B.12. *The set of hyperplanes can be colored with finitely many colors such that any two hyperplanes of the same color are disjoint.*

For a more precise result on the moduli space of cubic surfaces, see [ACT02a, Lemma 7.30, Theorem 7.32].

Proof. A very general argument, given in [Ber00, page 113] or [BHW11, Lemma 1.8], gives that if $H \in \mathcal{H}$, then $\text{Stab}_\Gamma(H)$ is a separable subgroup of Γ .

As in [DCS19, Proof of Theorem 1.1], ideas of Scott [Sco78] can then be adapted to find a finite Galois cover of $\Gamma \backslash B$ where the image of H does not intersect itself. (Some care is required in the non-uniform case: See [GR69, Theorem 1.2], or adapt arguments in [Ber00] to the complex hyperbolic case.) In the associated Galois cover the same holds for every hyperplane in ΓH . We can then take a common Galois cover of one such cover for every Γ orbit in $\mathcal{H}$. The result is a Galois cover of $\Gamma \backslash B$ where the image of every hyperplane in $\mathcal{H}$ does not intersect itself.

We then color each hyperplane by a set of colors equal to the images of elements of $\mathcal{H}$ in this cover. This gives the result. $\square$

Corollary B.13. *The set of stratum domains can be colored so any two of the same color have disjoint closures.*

Proof. One can color a stratum by the set of colors of the hyperplanes containing it. $\square$

Lemma B.14. *The set of cusp domains can be colored with finitely many colors such that any two of the same color are asynchronous.*

Proof. All the center domains can be assigned the same new color. Color each tree domain, viewed as a parallelism class of hyperplanes, with the set of colors of the hyperplanes it contains. $\square$

We can now conclude.

Lemma B.15. *$\mathfrak{S}$ is colorable.*

Proof. It suffices to handle each type of domain separately, and these are handled by the previous lemmas. $\square$

APPENDIX C. RELATIVE HIERARCHICAL HYPERBOLICITY (NOT REQUIRED)

For the definition of a relatively hierarchically hyperbolic space (RHHS) and group (RHHG), we recommend consulting [BHS17a, Definition 1.4, Section 1.2.3]; the original reference [BHS19, Definition 6.8] is similar but does not discuss the group version explicitly. The definitions of a RHHS and RHHG are identical to those of an HHS and HHG, except that the $\mathcal{C}(U)$ for U nest-minimal are now allowed to be arbitrary geodesic metric spaces. (The projection π_U should still be coarsely onto and coarsely Lipschitz.)

Minor modifications to our proof of [Theorem 1.2](#) prove that $\hat{\Gamma}$ is an RHHG even if [Assumption \(D\)](#) on cusps is dropped. [Appendix D](#) provides a number of examples where this is of interest.

Theorem C.1. *Suppose $\mathcal{H}$ and $\hat{\Gamma}$ are as in [Theorem 1.2](#) except that [Assumption \(D\)](#) is not assumed to hold. Then $\hat{\Gamma}$ is a relatively hierarchically hyperbolic group (RHHG).*

Almost all the proof proceeds almost verbatim the same. The main changes concern cusp domains. In a sense there may be fewer tree domains than before, and so the metric spaces $\mathcal{C}(U)$ attached to center domains U now must sometimes be larger. Their role expands enough that it becomes potentially misleading to continue to call them center domains, and so we will call them Heisenberg domains instead. For the purposes of this appendix, a center domain can thus be thought of as a special case of a Heisenberg domain that arises when [Assumption \(D\)](#) in fact does hold for the associated cusp.

We now comment in more detail on the sections requiring changes.

Section 3 and Appendix A: One can modify [Proposition 3.6](#) and [Proposition 3.7](#) simply by dropping the requirement that $\mathcal{C}(U)$ must be hyperbolic when U is nest-minimal, so they yield relative hierarchical hyperbolicity instead of hierarchical hyperbolicity. The proofs do not change; the only use of hyperbolicity as a tool in [Appendix A](#) is in [Lemma A.2](#), which only uses it for non-nest-minimal domains.

Section 6: It is natural to work with swaddling in a slightly more general context. Rather than working only with the Lie group $G = \mathbb{R}$, we now allow other connected Lie groups G with left invariant metrics. Maps that previously were $\mathbb{R}$ -equivariant are now required to be G -equivariant. We assume that, for every x, y, z , the G -equivariant map $m_{z,x} \circ m_{y,z} \circ m_{x,y}$ is multiplication by an element of a fixed subgroup of the center of G that is isomorphic, via a fixed isomorphism, to $\mathbb{R}$. [Lemmas 6.4](#) and [6.5](#) and its proofs are valid in this more general context, and we do not require any later results from [Section 6](#) in this more general context.

Remark C.2. The centrality assumption avoids having to work around the fact that a self-map of a G -torsor cannot be canonically identified with an element of G ; it distinguishes only a conjugacy class.

Section 7: This is the main section requiring changes, which begin in [Section 7.2](#). In comparison to [Lemma 7.3](#), we now have horosphere minus arrangement mod center of the form

$$\mathbb{C}^k \times \prod_{i=k+1}^{n-1} \mathbb{C}_i^\circ.$$

Unlike before, k can be positive, and even equal to $n - 1$ if the hyperplane arrangement doesn't go into the given horoball. The $\mathbb{C}^k$ does not factor canonically so is best thought of as an abstract vector space of dimension k . Our first observation is to determine the replacement for the center of the Heisenberg group.

Lemma C.3. *The connected component of the identity of the stabilizer of $\mathcal{H} \cap \partial T_c$ in the Heisenberg group H^{2n-1} is a Heisenberg group H^{2k+1} of dimension $2k + 1$.*

Note that when $k = 0$, this recovers the center of H^{2n-1} . The proof is left to the reader, and is implicit in [Section 7](#).

This H^{2k+1} will replace the center Z in most of our analysis. The space ∂T_c has the structure of a bundle

$$H^{2k+1} \rightarrow H^{2n-1} \rightarrow \prod_{i=k+1}^{n-1} \mathbb{C}_i.$$

It is this bundle and its variants that we use. The fibers are no longer real lines, but instead are torsors for H^{2k+1} . Given a path in $\prod_{i=k+1}^{n-1} \mathbb{C}_i$ we can lift it to a horizontal path in H^{2n-1} in two steps. First, lift it to a path in $\mathbb{C}^k \times \prod_{i=k+1}^{n-1} \mathbb{C}_i$, simply by keeping the $\mathbb{C}^k$ coordinate constant. Then, lift to a horizontal path in H^{2n-1} . This analysis shows that the monodromy for the bundle $H^{2k+1} \rightarrow H^{2n-1} \rightarrow \prod_{i=k+1}^{n-1} \mathbb{C}_i$ in fact lives in the center of H^{2n-1} , which is equal to the center of H^{2k+1} . The monodromy of a loop is again a sum of signed areas, exactly as in [Section 7.1](#).

[Section 7.3](#) requires only notational changes. We get that $\partial \hat{T}_c^\circ$ is the pullback of

$$\partial T_c^\circ \rightarrow \prod_{i=k+1}^{n-1} \mathbb{C}_i^\circ$$

to the universal cover $\prod_{i=k+1}^{n-1} \hat{\mathbb{C}}_i^\circ$ of its base. The discussion of tree domains in [Section 7.4](#) does not require modification, except to say there may now be fewer tree domains and to make straightforward modifications to [Lemma 7.3](#). We get one tree domain for each parallelism class of hyperplanes entering the horoball.

Remark C.4. If $k > 0$, then $\hat{\Gamma}$ is not an HHG. This follows from the Tits alternative for HHGs, which holds in a strong form: every finitely generated subgroup either is virtually abelian or contains a nonabelian free group [[DHS20](#), Theorem 4.1]. As in the proof of [Lemma 7.3](#), one can show directly that $\Gamma \cap H^{2k+1}$ is a cocompact lattice in H^{2k+1} , hence nilpotent but not virtually abelian.

The key changes are to [Section 7.5](#). Before we had $n - 1$ tree domains and one center domain per cusp. Now we have $n - 1 - k$ tree domains and one Heisenberg domain. If U is the Heisenberg domain, $\mathcal{C}(U)$ will still be a swaddling of $\partial \hat{T}_c \rightarrow \prod_{i=k+1}^{n-1} \hat{\mathbb{C}}_i$. The main difference is that now the swaddling is quasi-isometric not to a line but instead to H^{2k+1} . The proof of [Lemma 7.8](#) is essentially unchanged but must be rephrased because the connection one-form is now valued in the Lie algebra of H^{2k+1} ; compare for example to [[Ham17](#), Section 5.10].

In [Section 7.6](#), it is important that [Lemma 7.12](#) and its proof hold essentially unchanged.

Section 9: [Corollaries 9.8](#) and [9.10](#) and [Lemma 9.9](#) contain unnecessarily precise claims that now require straightforward modifications. For example, in [Corollary 9.8](#) it is still true that the maximal size of a chain is at most $n + 1$, but the (never used) criterion for equality need not hold if cusps need not contain 1-dimensional strata.

No other significant changes are required to the proof.

Remark C.5. $\hat{\Gamma}$ is only relatively hyperbolic in very special situations. Compare to [[BDM09](#), [Rus22](#), [ABR26](#), [BH13](#)]. See also the discussion after [[BDM09](#), Corollary 8.3] for prior results implying that most mapping class groups are not relatively hyperbolic.

Suppose $(G, \mathfrak{S})$ is an HHG for which the hyperbolic space of the nest-maximal domain is infinite diameter. It follows automatically that the Gromov boundary

of this hyperbolic space is non-empty; see, for example, [ABO19, Theorem 4.2]. Suppose also that there is a G -invariant subset $\mathfrak{S}_0 \subset \mathfrak{S}$, of size greater than 1, such that any two domains in $\mathfrak{S}_0$ can be joined by a sequence of domains in $\mathfrak{S}_0$, each orthogonal to the next, and such that the Gromov boundary of $\mathcal{C}(U)$ is non-empty for each $U \in \mathfrak{S}_0$. Then [ABR26, Theorem 6.12] gives that G is thick of order 1 and hence, by [BDM09, Corollary 7.9], G is not relatively hyperbolic. This can be applied to show that the $\hat{\Gamma}$ associated with the moduli space of cubic surfaces is not relatively hyperbolic. Here one should let $\mathfrak{S}_0$ be the set of all angle domains, and use the results mentioned in Remark 5.8.

APPENDIX D. OTHER EXAMPLES (NOT REQUIRED)

We give many examples to which Theorem 1.2 or Theorem C.1 applies. Mainly, we review a number of known constructions that produce arrangements of mutually orthogonal hyperplanes in the n -ball, invariant under lattices in $PU(n, 1)$.

D.1. The definition of orthogonality. Suppose V is a complex vector space and h is a hermitian form on V of signature $(m, 1)$. If v is a vector in V with $h(v, v) > 0$, we define the corresponding hyperplane to be

$$H_v = \{[w] \in \mathbb{P}V : h(w, v) = 0, h(w, w) < 0\}.$$

Note that the condition $h(v, v) > 0$ is exactly the condition required to make this a codimension 1 subspace of $\mathbb{C}\mathbb{H}^m$ (rather than being all of $\mathbb{C}\mathbb{H}^m$ or empty). We say v is positive when $h(v, v) > 0$. The first basic observation is the following.

Lemma D.1. *Suppose $r, s \in V$ are non-colinear positive vectors. Then $H_r \cap H_s \neq \emptyset$ if and only if*

$$|h(r, s)|^2 < h(r, r)h(s, s),$$

and they intersect orthogonally if and only if $h(r, s) = 0$.

For convenience we recall the proof of the first claim. Compare to the proof of [ACT02a, Lemma 7.28].

Proof. By definition, $H_r \cap H_s \neq \emptyset$ if and only if there is a negative vector orthogonal to both r and s . Such a vector exists if and only if the restriction of h to the span of r and s is positive definite, or in other words if and only if the matrix

$$\begin{pmatrix} h(r, r) & h(r, s) \\ h(s, r) & h(s, s) \end{pmatrix}$$

is positive definite. Since $h(r, r) > 0$, this is the case if and only if the determinant $h(r, r)h(s, s) - |h(r, s)|^2$ is positive. $\square$

D.2. A unit norm criterion for orthogonality. An arithmetic lattice “of the simplest type” in $PU(m, 1)$ can be defined as follows. Let $E \subset \mathbb{C}$ be a CM field, let $F \subset \mathbb{R}$ be its totally real subfield, and let $\mathcal{O}_E$ and $\mathcal{O}_F$ be the rings of integers of E and F . Let $V = E^{m+1}$, and let

$$h : V \times V \rightarrow E$$

be a hermitian form of signature $(m, 1)$. Each real embedding $\sigma : F \rightarrow \mathbb{R}$ gives a pair of complex conjugate maps $E \rightarrow \mathbb{C}$, and for all σ not the identity we assume

the corresponding pair of Galois conjugates of h have signature $(m+1, 0)$. With these assumptions,

$$PU(\mathcal{O}_E, h) \subset PU(V \otimes \mathbb{C}, h)$$

is a lattice, and it is exactly lattices commensurable with such $PU(\mathcal{O}_E, h)$ which are called “of the simplest type”.

Remark D.2. This lattice is cocompact if $F \neq \mathbb{Q}$.

Note that arithmetic lattices not of the simplest type involve non-commutative central simple division algebras with involutions of the second kind. See for example the brief discussion in [LIP24, Section 3.2], or the discussion in [ES14, ST22] as well as earlier sources.

In this common setup we have the following orthogonality criterion.

Lemma D.3. *Suppose non-colinear positive vectors $r, s \in V$ have that $h(r, r)$ and $h(s, s)$ are units in $\mathcal{O}_F$ and $h(r, s)$ is in $\mathcal{O}_E$. Then H_r and H_s are orthogonal or disjoint.*

The main case to keep in mind is $h(r, r) = 1 = h(s, s)$.

Proof. Note that $h(r, s)h(s, r) = |h(r, s)|^2$ is in F . Suppose H_r and H_s intersect. Then Lemma D.1 gives

$$|h(r, s)|^2 < h(r, r)h(s, s).$$

For non-trivial Galois conjugates σ as above, the fact that the Galois conjugate forms are definite, together with Cauchy-Schwarz, gives

$$\sigma(h(r, s)h(s, r)) \leq \sigma(h(r, r))\sigma(h(s, s)).$$

Appropriately multiplying all these together gives

$$N_{F/\mathbb{Q}}(h(r, s)h(s, r)) < N_{F/\mathbb{Q}}(h(r, r))N_{F/\mathbb{Q}}(h(s, s)).$$

As a product of norms of units, the right side is ± 1 . The left side is non-negative, since it is a product of non-negative terms

$$\sigma(h(r, s)h(s, r)) = \sigma(h(r, s))\overline{\sigma(h(r, s))}.$$

Therefore $N_{F/\mathbb{Q}}(|h(r, s)|^2) = 0$, and hence $h(r, s) = 0$ as desired. $\square$

Already this can give infinitely many examples where Theorem 1.2 or Theorem C.1 applies. For example, let E be a cyclotomic field, let F be its totally real subfield, let $a \in \mathcal{O}_F$ be a negative number all of whose non-trivial conjugates are positive, and define

$$h(z, w) = z_1\bar{w}_1 + \cdots + z_m\bar{w}_m + az_{m+1}\bar{w}_{m+1}.$$

Then Lemma D.3 gives that the unit norm elements in $\mathcal{O}_E^{m+1} \subset E^{m+1}$ define an orthogonal arrangement invariant under a lattice. Whether Theorem 1.2 applies depends on whether its hypothesis (D) holds. This is automatic when the lattice is cocompact, for example whenever $F \neq \mathbb{Q}$. Compare to [dGF23, Example 2.13].

In the non-cocompact setting, it is difficult in general to get that every cusp contains a 1-dimensional stratum. Take for example the case when $\mathcal{O}_E$ is the ring of Eisenstein integers, and consider the standard hermitian form

$$h(z, w) = z_1\bar{w}_1 + \cdots + z_m\bar{w}_m - z_{m+1}\bar{w}_{m+1}.$$

The lattice here is $PU(\mathcal{O}_E, h)$ as above. In general, cusps correspond to orbits of primitive isotropic lines, and to each such line ℓ one can associate the positive

definite lattice $\ell^\perp/\ell$. In this specific case this gives a bijection between cusps and positive definite hermitian self-dual lattices of rank $m - 1$, and there is only one such when $m \leq 6$. See [All00b, Section 7] for this correspondence, and [Fei78] for the original classification of such lattices. Thus there is only one cusp when $m \leq 6$, and it contains a 1-dimensional stratum, but already for $m = 7$ there is a second cusp that is disjoint from the hyperplane arrangement.

Consulting [AS08, Remark 1.2, Examples 1.1, Proposition 3.2, Table 2, Corollary 5.3] and using [All00b, Theorem 7.1], it is possible to see that the moduli spaces

$$M_{3,0}, M_{4,1}, M_{5,2}, M_{6,3}, M_{7,4}, M_{8,5}$$

of K3 surfaces with non-symplectic automorphisms of order 3 with no fixed curve of higher genus are all quotients of complex balls with hyperplane arrangements defined by the norm 1 elements with entries in the ring of Eisenstein integers. In these cases the balls have dimensions 6, 5, 4, 3, 2, 1 respectively, and the hyperplane arrangements are orthogonal by Lemma D.3. See [MOT15] for more on these spaces, and note in particular that $M_{5,2}$ is isomorphic to the moduli space of smooth complex cubic surfaces. The previous paragraph shows that each corresponding ball quotient has a single cusp, and that it lies in the closure of a dimension 1 stratum.

D.3. Moduli spaces of abelian varieties with extra structure (unitary Shimura varieties of PEL type). Take E and F as in Appendix D.2, and assume given a free $\mathcal{O}_E$ -module Λ of rank $m + 1$ and a hermitian form

$$h : \Lambda \times \Lambda \rightarrow \mathcal{O}_E.$$

As in Appendix D.2, with $V = \Lambda \otimes_{\mathcal{O}_E} E$, we assume h has signature $(m, 1)$ and its other Galois conjugates have signature $(m + 1, 0)$. We define a short root to be any $r \in \Lambda$ with $h(r, r) = 1$. Lemma D.3 immediately implies that the hyperplane arrangement defined by the short roots is orthogonal, and so Theorem C.1 (and sometimes Theorem 1.2) applies.

This setup includes in particular the class of “admissible hermitian lattices” of [dGF23, Section 2.1], which additionally places a restriction on the different ideal of $\mathcal{O}_E$. In this setting, [dGF23, Section 11] shows that the ball quotient associated to (Λ, h) is a moduli space of polarized $\mathcal{O}_E$ -linear abelian varieties, and the divisor defined by the short roots has a natural modular interpretation as the locus where an additional suitable morphism from a fixed CM torus exists. In this context the divisor is also known as the norm-one Kudla–Rapoport divisor.

The case of quadratic imaginary fields E is especially nice, and the condition on different ideals is automatic [dGF23, Example 2.12]. In that case the ball quotient minus the norm 1 divisor parametrizes abelian varieties with a suitable action of $\mathcal{O}_E$, of the fixed signature and polarization type, but without a factor of the fixed CM elliptic curve $\mathbb{C}/\mathcal{O}_E$ as a polarized $\mathcal{O}_E$ summand.

D.4. Orthogonality via root lattices. We now consider the following setup. Suppose Λ is a lattice over $\mathbb{Z}$ of signature $(N, 2)$, and $g \in O(\Lambda)$ has order $d > 1$. Let ζ be a primitive d -th root of unity, and assume the eigenspace

$$V = \ker(g - \zeta) \subset \Lambda \otimes_{\mathbb{Z}} \mathbb{C}$$

has signature $(m, 1)$ for some m , with respect to the hermitian form $h(v, w) = (v, \overline{w})$. Then we can consider the complex hyperbolic space $\mathbb{CH}^m$ defined by V , with the

hyperplane arrangement defined by the (g -orbits of) norm 2 elements of Λ that define hyperplanes in $\mathbb{CH}^m$.

Proposition D.4. *If d is co-prime to 30, and g does not have any non-zero fixed vectors, then this arrangement is orthogonal.*

The assumption could be weakened, but here we prefer a simple statement. The proof can be compared to [DvGK05, Section 8] and is omitted. The key is to show the following: Let Q be an irreducible ADE root lattice with an isometry u of order k co-prime to 30 and with no non-zero fixed vectors. Then Q has type A_{k-1} , the isometry u is a k -cycle in its Weyl group S_k , and in particular every non-trivial eigenspace of u on $Q \otimes \mathbb{C}$ is one dimensional.

Remark D.5. The case $d = 7$ is the motivating example for this subsection. The moduli space of K3 surfaces with a non-symplectic automorphism of order 7 has two components in [AST11, Theorem 9.5], both of which are quotients of two-dimensional balls minus hyperplane arrangements which are orthogonal by Proposition D.4. The discussion at the beginning of [AST11, Section 9] gives the setup. In this case there are no cusps by Remark D.2, so Theorem 1.2 applies.

D.5. Orthogonality via commuting isometries. We start with the following well known observation, whose proof is left to the reader. Compare to [Der24, Remark 2.3].

Lemma D.6. *Suppose $\Gamma_0 \backslash \mathbb{CH}^m$ is a manifold and H is a finite group of holomorphic isometries of $\Gamma_0 \backslash \mathbb{CH}^m$ such that for all $x \in \Gamma_0 \backslash \mathbb{CH}^m$ we have that the stabilizer of x in H is Abelian. Let $\mathcal{D} \subset \Gamma_0 \backslash \mathbb{CH}^m$ be the union of the complex codimension 1 components of the fixed point sets of elements of H . Then the preimage $\mathcal{H}$ of $\mathcal{D}$ in $\mathbb{CH}^m$ is an orthogonal arrangement.*

There are a few known line arrangements in $\mathbb{CP}^2$ with finite (branched) Galois covers to which this applies. The complements of these line arrangements in $\mathbb{CP}^2$ have the form $\Gamma \backslash (\mathbb{CH}^2 - \mathcal{H})$ with $\mathcal{H}$ orthogonal by Lemma D.6. See [Hir83, BHH87, Tre16]. Important examples include Klein’s arrangement, studied for example in [Nar90], and the Hesse arrangement, studied for example in [Kan23]. In both cases there is in fact more than one complex hyperbolic metric known on the complement of the line arrangement, but in both cases there is (at least) one with Γ_0 cocompact where Lemma D.6 applies on a finite cover. See [Tre16, Definition 3.1, Tables 5.2 and 5.4, Theorem 6.4] as one possible starting point; and see [Der19, CHL05] for some additional context. So we get that Theorem 1.2 applies to the fundamental groups of complements in $\mathbb{CP}^2$ of the Hesse and Klein arrangements.

D.6. Additional examples and special cases. Here are some notable examples.

Stable cubic threefolds: A cubic threefold means a hypersurface in $\mathbb{CP}^4$ defined by a cubic equation. Stability is a concept from Geometric Invariant Theory, that turns out to be equivalent to every singularity having one of the types A_1, A_2, A_3, A_4 [All03, Yok02]. The moduli space of stable cubic threefolds, or equivalently the moduli space of K -stable cubic threefolds [LX19], may be identified with the complex 10-ball, minus a family of disjoint hyperplanes, modulo a discrete group [ACT11, Theorems 3.7, 7.2], [LS07]. Assumption (D) of Theorem 1.2 fails since there are no 1-dimensional strata and the complex hyperbolic orbifold is not compact (see [ACT11, Theorem 4.10]), but Theorem C.1 applies. Caution: this

identification with a ball quotient is an isomorphism of varieties, but not one of orbifolds. (Away from additional hyperplanes, that represent singular threefolds and form a non-orthogonal arrangement, it is also an orbifold isomorphism.)

Stable quartic plane curves: This is similar to the previous example, with the 6-ball in place of the 10-ball [Kon00]. Now stability is equivalent to having only A_1 or A_2 singularities. The same caution applies, regarding orbifold structure along the hyperplanes representing singular curves. We remark that smooth quartic plane curves, up to projective equivalence, are the same as non-hyperelliptic genus 3 curves, up to isomorphism. The points of B^6 , whose removal leaves the stable locus, represent hyperelliptic genus 3 curves. So, by adding them back, the full ball quotient is the moduli space of genus 3 curves with mild degenerations allowed.

Moduli of 5 points, or smooth complex binary quintics: Here the moduli space, also known as $\mathcal{M}_{0,5}/S_5$, can be realized as the complement of an orthogonal arrangement in a compact ball quotient of dimension 2, so Theorem 1.2 applies. This has been known in some form for a long time [Shi64, Shi63, DM86], but we recommend [dGF25, Section 4] as a reference that treats orthogonality explicitly. Up to finite covers, this moduli space is also the moduli space of smooth degree 4 del Pezzo surfaces [HKT16, Section 6], [Dol16b], and also a certain moduli space of K3 surfaces with order 5 non-symplectic automorphisms [Kon07, AST11].

Remark D.7. We also have that

$$\mathcal{M}_{0,5} \simeq (\mathbb{CP}^1 - \{0, 1, \infty\})^2 - \Delta \simeq \Gamma(2) \times \Gamma(2) \backslash (\mathbb{H}^2 \times \mathbb{H}^2 - \mathcal{H}),$$

where Δ is the diagonal and $\mathcal{H}$ is its preimage and $\Gamma(2) \subset PSL(2, \mathbb{Z})$ is the principal level 2 congruence subgroup. Since $\mathbb{H}^2 \times \mathbb{H}^2$ is the $SO(2, n)$ symmetric space for $n = 2$, this shows that $\mathcal{M}_{0,5}$ fits both within the complex hyperbolic and type IV framework.

Moduli of 6 points, or smooth complex binary sextics: Here the moduli space, also known as $\mathcal{M}_{0,6}/S_6$, is the three-dimensional Eisenstein example presented above, so Theorem 1.2 applies. Again this goes back at least to [DM86], but see also [ACT07, Theorem 3] and [DK07, Section 12] as a reference.

Deligne-Mostow-Thurston examples and mapping class groups: Deligne-Mostow and later Thurston gave a number of examples of ball quotients [DM86, Thu98]. In the Deligne-Mostow point of view, they parameterize certain cyclic covers of the sphere, and in the Thurston point of view they parametrize certain cone metrics on the sphere. These ball quotients come equipped with a natural divisor corresponding to certain collisions between points. (Other collisions between points live “at infinity” on the ball quotient, or are not allowed.)

We now refer to the table of 94 examples in [Thu98, Appendix]. In each, the divisor is orthogonal if, for every triple of numerators, their sum is at least as large as the sum of the remaining numerators. The entries on Thurston’s list with this property are: 1, 2, 8, 9, 42, 43, 47, 48, 72, 73, 74, 75, 77, 78, 83, 84, 87, 89, 94. Each is a lattice in $PU(2, 1)$, except the first entry, which is a lattice in $PU(3, 1)$. All are cocompact except 1, 2, 8, 42, 43, 73, and Assumption (D) holds in all these cases. Thus, Theorem 1.2 applies to all 19 examples listed above.

The moduli of cubic surfaces is not a Deligne-Mostow-Thurston example [All00b, Theorem 8.4] but it is commensurable with one [Dor04, Theorem 3], [DvGK05,

Section 3], [CMGH25, Section 3]. This Deligne-Mostow-Thurston example is not one of the orthogonal ones listed above, and the commensurability result involves passing to a cover and removing only certain components of the divisor discussed above.

We will not make any attempt to exhaustively list the coincidences and relationships among the examples we have presented, but it should be noted that examples number 1 and 9 on Thurston’s list correspond to the cases of 6 and 5 points respectively on $\mathbb{CP}^1$ discussed above, and we wish to emphasize that the full mapping class groups for the 5 and 6 times punctured sphere are valid examples of $\hat{\Gamma}$ to which Theorem 1.2 applies. More generally, the other orthogonal examples on Thurston’s list show that various partially labeled finite index subgroups of these two mapping class groups can be realized in different ways as a $\hat{\Gamma}$ to which Theorem 1.2 applies.

Remark D.8. Our groups are similar enough to mapping class groups that we feel compelled to briefly sketch why, in the cubic surface case, $\hat{\Gamma}$ is not abstractly commensurable with a mapping class group. By [BHS21, Theorem 1.15] and Corollary 9.10, if this $\hat{\Gamma}$ were abstractly commensurable with a mapping class group of a genus g surface with k punctures, we’d have $3g - 3 + k = 4$. For such surfaces, the abstract commensurator is the extended mapping class group; see [Iva97, Kor99] and [BM06, Theorem 6]. On the other hand, the Fermat cubic gives a finite subgroup $(\mathbb{Z}/3)^3 \rtimes S_4 < \hat{\Gamma}$ of order 648 [Dol12, Theorem 9.5.8], and this subgroup injects into the abstract commensurator of $\hat{\Gamma}$. But finite subgroups of the extended mapping class group have order at most 60 in the cases $S_{0,7}$, $S_{1,4}$, and $S_{2,1}$ [May77, Ker83], so we can conclude that $\hat{\Gamma}$ is not abstractly commensurable with a mapping class group. This argument can be compared to arguments in [CLM14, BM19, Sor21], and is based on a conversation with ChatGPT 5.5 Pro.

Additionally invoking quasi-isometric rigidity of mapping class groups [Ham07, BKMM12, BHS21], it is possible to prove that the $\hat{\Gamma}$ in the cubic surface case is not even quasi-isometric to a mapping class group.

REFERENCES

- [ABD21] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham, *Largest acylindrical actions and stability in hierarchically hyperbolic groups*, Trans. Amer. Math. Soc. Ser. B **8** (2021), 66–104, With an appendix by Daniel Berlyne and Jacob Russell.
- [ABL26] Daniel Allcock, Tathagata Basak, and Eduard Looijenga, *The moduli space of cubic surfaces and its orbifold fundamental group*, 2026, preprint.
- [ABO19] Carolyn Abbott, Sahana H. Balasubramanya, and Denis Osin, *Hyperbolic structures on groups*, Algebr. Geom. Topol. **19** (2019), no. 4, 1747–1835.
- [ABR25] Carolyn Abbott, Jason Behrstock, and Jacob Russell, *Structure invariant properties of the hierarchically hyperbolic boundary*, J. Topol. Anal. **17** (2025), no. 3, 859–903.
- [ABR26] ———, *Relative Hyperbolicity, Thickness, and the Hierarchically Hyperbolic Boundary*, Michigan Math. J. **76** (2026), no. 3, 481–526.
- [ACT02a] Daniel Allcock, James A. Carlson, and Domingo Toledo, *The complex hyperbolic geometry of the moduli space of cubic surfaces*, J. Algebraic Geom. **11** (2002), no. 4, 659–724.
- [ACT02b] ———, *Orthogonal complex hyperbolic arrangements*, Symposium in Honor of C. H. Clemens (Salt Lake City, UT, 2000), Contemp. Math., vol. 312, Amer. Math. Soc., Providence, RI, 2002, pp. 1–8.
- [ACT07] ———, *Hyperbolic geometry and the moduli space of real binary sextics*, Arithmetic and geometry around hypergeometric functions, Progr. Math., vol. 260, Birkhäuser, Basel, 2007, pp. 1–22.

- [ACT11] ———, *The moduli space of cubic threefolds as a ball quotient*, Mem. Amer. Math. Soc. **209** (2011), no. 985, xii+70.
- [AE25] Valery Alexeev and Philip Engel, *On lattice-polarized K3 surfaces*, 2025, arXiv:2505.22557.
- [AEH24] Valery Alexeev, Philip Engel, and Changho Han, *Compact moduli of K3 surfaces with a nonsymplectic automorphism*, Trans. Amer. Math. Soc. Ser. B **11** (2024), 144–163.
- [AF02] Daniel Allcock and Eberhard Freitag, *Cubic surfaces and Borcherds products*, Comment. Math. Helv. **77** (2002), no. 2, 270–296.
- [AH26] P  n  lope Azuelos and Mark Hagen, *Periodic quasiflats in hierarchically hyperbolic spaces*, 2026, arXiv:2608.01513.
- [AHW25] Francisco Arana-Herrera and Alex Wright, *The geometry of totally geodesic subvarieties of moduli spaces of Riemann surfaces*, 2025, arXiv:2412.16330.
- [AKP24] Stephanie Alexander, Vitali Kapovitch, and Anton Petrunin, *Alexandrov geometry—foundations*, Graduate Studies in Mathematics, vol. 236, American Mathematical Society, Providence, RI, [2024]   2024.
- [All98] D. Allcock, *An isoperimetric inequality for the Heisenberg groups*, Geom. Funct. Anal. **8** (1998), no. 2, 219–233.
- [All00a] Daniel Allcock, *Asphericity of moduli spaces via curvature*, J. Differential Geom. **55** (2000), no. 3, 441–451.
- [All00b] ———, *New complex- and quaternion-hyperbolic reflection groups*, Duke Math. J. **103** (2000), no. 2, 303–333.
- [All00c] ———, *The period lattice for Enriques surfaces*, Math. Ann. **317** (2000), no. 3, 483–488.
- [All03] ———, *The moduli space of cubic threefolds*, J. Algebraic Geom. **12** (2003), no. 2, 201–223.
- [All09] ———, *A monstrous proposal*, Groups and symmetries, CRM Proc. Lecture Notes, vol. 47, Amer. Math. Soc., Providence, RI, 2009, pp. 17–24.
- [All13] ———, *Completions, branched covers, Artin groups, and singularity theory*, Duke Math. J. **162** (2013), no. 14, 2645–2689.
- [AS08] Michela Artebani and Alessandra Sarti, *Non-symplectic automorphisms of order 3 on K3 surfaces*, Math. Ann. **342** (2008), no. 4, 903–921.
- [AST11] Michela Artebani, Alessandra Sarti, and Shingo Taki, *K3 surfaces with non-symplectic automorphisms of prime order*, Math. Z. **268** (2011), no. 1-2, 507–533, With an appendix by Shigeyuki Kond  .
- [Aus60] Louis Auslander, *Bieberbach’s theorems on space groups and discrete uniform subgroups of Lie groups*, Ann. of Math. (2) **71** (1960), 579–590.
- [Bar18] Arthur Bartels, *K-theory and actions on Euclidean retracts*, Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures, World Sci. Publ., Hackensack, NJ, 2018, pp. 1041–1062.
- [BB19] Arthur Bartels and Mladen Bestvina, *The Farrell-Jones conjecture for mapping class groups*, Invent. Math. **215** (2019), no. 2, 651–712.
- [BDM09] Jason Behrstock, Cornelia Dru  u, and Lee Mosher, *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*, Math. Ann. **344** (2009), no. 3, 543–595.
- [Bel12a] Igor Belegradek, *Complex hyperbolic hyperplane complements*, Math. Ann. **353** (2012), no. 2, 545–579.
- [Bel12b] ———, *Rigidity and relative hyperbolicity of real hyperbolic hyperplane complements*, Pure Appl. Math. Q. **8** (2012), no. 1, 15–51.
- [Ber00] Nicolas Bergeron, *Premier nombre de Betti et spectre du laplacien de certaines vari  t  s hyperboliques*, Enseign. Math. (2) **46** (2000), no. 1-2, 109–137.
- [BFJS26] Gregorio Baldi, Benson Farb, Ariyan Javanpeykar, and Matthew Stover, *Automorphisms of the moduli space of smooth cubic surfaces and its fundamental group*, 2026, arXiv:2605.16658.
- [BH99] Martin R. Bridson and Andr   Haeffliger, *Metric spaces of non-positive curvature*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999.

- [BH13] Igor Belegradek and G. Christopher Hruska, *Hyperplane arrangements in negatively curved manifolds and relative hyperbolicity*, Groups Geom. Dyn. **7** (2013), no. 1, 13–38.
- [BHH87] Gottfried Barthel, Friedrich Hirzebruch, and Thomas Höfer, *Geradenkonfigurationen und Algebraische Flächen*, Aspects of Mathematics, D4, Friedr. Vieweg & Sohn, Braunschweig, 1987.
- [BHMS24] Jason Behrstock, Mark Hagen, Alexandre Martin, and Alessandro Sisto, *A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups*, J. Topol. **17** (2024), no. 3, Paper No. e12351, 94.
- [BHS17a] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups*, Proc. Lond. Math. Soc. (3) **114** (2017), no. 5, 890–926.
- [BHS17b] ———, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. **21** (2017), no. 3, 1731–1804.
- [BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula*, Pacific J. Math. **299** (2019), no. 2, 257–338.
- [BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Quasiflats in hierarchically hyperbolic spaces*, Duke Math. J. **170** (2021), no. 5, 909–996.
- [BHW11] Nicolas Bergeron, Frédéric Haglund, and Daniel T. Wise, *Hyperplane sections in arithmetic hyperbolic manifolds*, J. Lond. Math. Soc. (2) **83** (2011), no. 2, 431–448.
- [BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher, *Geometry and rigidity of mapping class groups*, Geom. Topol. **16** (2012), no. 2, 781–888.
- [BLW14] Matthew Strom Borman, Tian-Jun Li, and Weiwei Wu, *Spherical Lagrangians via ball packings and symplectic cutting*, Selecta Math. (N.S.) **20** (2014), no. 1, 261–283.
- [BM06] Jason Behrstock and Dan Margalit, *Curve complexes and finite index subgroups of mapping class groups*, Geom. Dedicata **118** (2006), 71–85.
- [BM19] Tara E. Brendle and Dan Margalit, *Normal subgroups of mapping class groups and the metaconjecture of Ivanov*, J. Amer. Math. Soc. **32** (2019), no. 4, 1009–1070.
- [Bow12] B. H. Bowditch, *Relatively hyperbolic groups*, Internat. J. Algebra Comput. **22** (2012), no. 3, 1250016, 66.
- [CHL05] Wim Couwenberg, Gert Heckman, and Eduard Looijenga, *Geometric structures on the complement of a projective arrangement*, Publ. Math. Inst. Hautes Études Sci. (2005), no. 101, 69–161.
- [CHRT25] Saajid Chowdhury, Hechen Hu, Matthew Romney, and Adam Tsou, *On $\text{CAT}(\kappa)$ surfaces*, Differential Geom. Appl. **101** (2025), Paper No. 102307, 21.
- [CLM14] Matt Clay, Christopher J. Leininger, and Dan Margalit, *Abstract commensurators of right-angled Artin groups and mapping class groups*, Math. Res. Lett. **21** (2014), no. 3, 461–467.
- [CMGH25] Sebastian Casalaina-Martin, Samuel Grushevsky, and Klaus Hulek, *The birational geometry of moduli of cubic surfaces and cubic surfaces with a line*, Moduli **2** (2025), Paper No. e1, 24.
- [CMGHL24] Sebastian Casalaina-Martin, Samuel Grushevsky, Klaus Hulek, and Radu Laza, *Non-isomorphic smooth compactifications of the moduli space of cubic surfaces*, Nagoya Math. J. **254** (2024), 315–365.
- [DCS19] Luca F. Di Cerbo and Matthew Stover, *Punctured spheres in complex hyperbolic surfaces and bielliptic ball quotient compactifications*, Trans. Amer. Math. Soc. **372** (2019), no. 7, 4627–4646.
- [DDL24] Spencer Dowdall, Matthew G. Durham, Christopher J. Leininger, and Alessandro Sisto, *Extensions of Veech groups II: Hierarchical hyperbolicity and quasi-isometric rigidity*, Comment. Math. Helv. **99** (2024), no. 1, 149–228.
- [Der19] Martin Deraux, *Non-arithmetic lattices and the Klein quartic*, J. Reine Angew. Math. **754** (2019), 253–279.
- [Der24] ———, *Mirror stabilizers for lattice complex hyperbolic triangle groups*, Geom. Dedicata **218** (2024), no. 3, Paper No. 58, 29.
- [DG] François Dahmani and Vincent Guirardel, *Rotating families, and the structure of some normal subgroups in groups acting on hyperbolic spaces*, <https://perso.univ-rennes1.fr/vincent.guirardel/papiers/rotating.pdf>.

- [dGF23] Olivier de Gaay Fortman, *Non-arithmetic hyperbolic orbifolds attached to unitary Shimura varieties*, 2023, arXiv:2301.01598.
- [dGF25] Olivier de Gaay Fortman, *Hyperbolic geometry and real moduli of five points on the line*, *Moduli* **2** (2025), Paper No. e20, 29.
- [DHS17] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto, *Boundaries and automorphisms of hierarchically hyperbolic spaces*, *Geom. Topol.* **21** (2017), no. 6, 3659–3758.
- [DHS20] ———, *Correction to the article Boundaries and automorphisms of hierarchically hyperbolic spaces*, *Geom. Topol.* **24** (2020), no. 2, 1051–1073.
- [DK07] Igor V. Dolgachev and Shigeyuki Kondō, *Moduli of K3 surfaces and complex ball quotients*, *Arithmetic and geometry around hypergeometric functions*, *Progr. Math.*, vol. 260, Birkhäuser, Basel, 2007, pp. 43–100.
- [DM86] P. Deligne and G. D. Mostow, *Monodromy of hypergeometric functions and non-lattice integral monodromy*, *Inst. Hautes Études Sci. Publ. Math.* (1986), no. 63, 5–89.
- [DMS23] Matthew G. Durham, Yair N. Minsky, and Alessandro Sisto, *Stable cubulations, bicomings, and barycenters*, *Geom. Topol.* **27** (2023), no. 6, 2383–2478.
- [DMS25] Matthew Gentry Durham, Yair Minsky, and Alessandro Sisto, *Asymptotically CAT(0) metrics, Z-structures, and the Farrell-Jones conjecture*, 2025, arXiv:2504.17048.
- [Dol96] I. V. Dolgachev, *Mirror symmetry for lattice polarized K3 surfaces*, *J. Math. Sci* **81** (1996), no. 3, 2599–2630.
- [Dol12] Igor V. Dolgachev, *Classical algebraic geometry*, Cambridge University Press, Cambridge, 2012, A modern view.
- [Dol16a] ———, *A brief introduction to Enriques surfaces*, *Development of moduli theory—Kyoto 2013*, *Adv. Stud. Pure Math.*, vol. 69, Math. Soc. Japan, [Tokyo], 2016, pp. 1–32.
- [Dol16b] ———, *Rational self-maps of moduli spaces*, *Pure Appl. Math. Q.* **12** (2016), no. 3, 335–351.
- [Dor04] Brent R. Doran, *Moduli Space of Cubic Surfaces as Ball Quotient via Hypergeometric Functions*, 2004, arXiv:0404062.
- [DSU17] Tushar Das, David Simmons, and Mariusz Urbański, *Geometry and dynamics in Gromov hyperbolic metric spaces*, *Mathematical Surveys and Monographs*, vol. 218, American Mathematical Society, Providence, RI, 2017, With an emphasis on non-proper settings.
- [DvGK05] I. Dolgachev, B. van Geemen, and S. Kondō, *A complex ball uniformization of the moduli space of cubic surfaces via periods of K3 surfaces*, *J. Reine Angew. Math.* **588** (2005), 99–148.
- [ES14] Vincent Emery and Matthew Stover, *Covolumes of nonuniform lattices in $PU(n, 1)$* , *Amer. J. Math.* **136** (2014), no. 1, 143–164.
- [Eva10] J. D. Evans, *Lagrangian spheres in del Pezzo surfaces*, *J. Topol.* **3** (2010), no. 1, 181–227.
- [Eva11] Jonathan David Evans, *Symplectic mapping class groups of some Stein and rational surfaces*, *J. Symplectic Geom.* **9** (2011), no. 1, 45–82.
- [Far98] B. Farb, *Relatively hyperbolic groups*, *Geom. Funct. Anal.* **8** (1998), no. 5, 810–840.
- [Fei78] Walter Feit, *Some lattices over $\mathbb{Q}(\sqrt{-3})$* , *J. Algebra* **52** (1978), no. 1, 248–263.
- [FFMS26] Francesco Fournier-Facio, Giorgio Mangioni, and Alessandro Sisto, *Bounded cohomology, quotient extensions, and hierarchical hyperbolicity*, *J. Topol. Anal.* (2026), 1–34.
- [FJ93] F. T. Farrell and L. E. Jones, *Isomorphism conjectures in algebraic K-theory*, *J. Amer. Math. Soc.* **6** (1993), no. 2, 249–297.
- [FS23] Roberto Frigerio and Alessandro Sisto, *Central extensions and bounded cohomology*, *Ann. H. Lebesgue* **6** (2023), 225–258.
- [FSW25] Hanlong Fang, Luca Schaffler, and Xian Wu, *Fineness and smoothness of a KSBA moduli of marked cubic surfaces*, *Proc. Amer. Math. Soc.* **153** (2025), no. 12, 5133–5146.
- [Ger92] S. M. Gersten, *Bounded cocycles and combings of groups*, *Internat. J. Algebra Comput.* **2** (1992), no. 3, 307–326.

- [GGH⁺26] David Gabai, David Gay, Daniel Hartman, Vyacheslav Krushkal, and Mark Powell, *Pseudo-isotopies of simply connected 4-manifolds*, Forum Math. Pi **14** (2026), Paper No. e9, 25.
- [GHM⁺24] Daniel Groves, Peter Haïssinsky, Jason F. Manning, Damian Osajda, Alessandro Sisto, and Genevieve S. Walsh, *Drilling hyperbolic groups*, 2024, arXiv:2406.14667.
- [GHS13] V. Gritsenko, K. Hulek, and G. K. Sankaran, *Moduli of K3 surfaces and irreducible symplectic manifolds*, Handbook of moduli. Vol. I, Adv. Lect. Math. (ALM), vol. 24, Int. Press, Somerville, MA, 2013, pp. 459–526.
- [GKS21] Patricio Gallardo, Matt Kerr, and Luca Schaffler, *Geometric interpretation of toroidal compactifications of moduli of points in the line and cubic surfaces*, Adv. Math. **381** (2021), Paper No. 107632, 48.
- [Gol99] William M. Goldman, *Complex hyperbolic geometry*, Oxford Mathematical Monographs, The Clarendon Press, Oxford University Press, New York, 1999, Oxford Science Publications.
- [GR69] H. Garland and M. S. Raghunathan, *Fundamental domains for lattices in rank one semisimple Lie groups*, Proc. Nat. Acad. Sci. U.S.A. **62** (1969), 309–313.
- [GR13] Giovanni Gandini and Henrik Rüping, *The Farrell-Jones conjecture for graph products*, Algebr. Geom. Topol. **13** (2013), no. 6, 3651–3660.
- [Hag23] Mark Hagen, *Non-colorable hierarchically hyperbolic groups*, Internat. J. Algebra Comput. **33** (2023), no. 2, 337–350.
- [Ham07] Ursula Hamenstaedt, *Geometry of the mapping class groups III: Quasi-isometric rigidity*, 2007, arXiv:0512429.
- [Ham17] Mark J. D. Hamilton, *Mathematical gauge theory*, Universitext, Springer, Cham, 2017, With applications to the standard model of particle physics.
- [Har79] Joe Harris, *Galois groups of enumerative problems*, Duke Math. J. **46** (1979), no. 4, 685–724.
- [HHP23] Thomas Haettel, Nima Hoda, and Harry Petyt, *Coarse injectivity, hierarchical hyperbolicity and semihyperbolicity*, Geom. Topol. **27** (2023), no. 4, 1587–1633.
- [Hir83] F. Hirzebruch, *Arrangements of lines and algebraic surfaces*, Arithmetic and geometry, Vol. II, Progr. Math., vol. 36, Birkhäuser, Boston, MA, 1983, pp. 113–140.
- [HKT09] Paul Hacking, Sean Keel, and Jenia Tevelev, *Stable pair, tropical, and log canonical compactifications of moduli spaces of del Pezzo surfaces*, Invent. Math. **178** (2009), no. 1, 173–227.
- [HKT16] Brendan Hassett, Andrew Kresch, and Yuri Tschinkel, *On the moduli of degree 4 del Pezzo surfaces*, Development of moduli theory—Kyoto 2013, Adv. Stud. Pure Math., vol. 69, Math. Soc. Japan, [Tokyo], 2016, pp. 349–386.
- [HL02] Gert Heckman and Eduard Looijenga, *The moduli space of rational elliptic surfaces*, Algebraic geometry 2000, Azumino (Hotaka), Adv. Stud. Pure Math., vol. 36, Math. Soc. Japan, Tokyo, 2002, pp. 185–248.
- [HMS24] Mark Hagen, Alexandre Martin, and Alessandro Sisto, *Extra-large type Artin groups are hierarchically hyperbolic*, Math. Ann. **388** (2024), no. 1, 867–938.
- [HMS26] Mark Hagen, Giorgio Mangioni, and Alessandro Sisto, *A combinatorial structure for many hierarchically hyperbolic spaces*, Pacific J. Math. **341** (2026), no. 2, 305–377.
- [Hor78a] Eiji Horikawa, *On the periods of Enriques surfaces. I*, Math. Ann. **234** (1978), no. 1, 73–88.
- [Hor78b] ———, *On the periods of Enriques surfaces. II*, Math. Ann. **235** (1978), no. 3, 217–246.
- [HRSS25] Mark Hagen, Jacob Russell, Alessandro Sisto, and Davide Spriano, *Equivariant hierarchically hyperbolic structures for 3-manifold groups via quasimorphisms*, Ann. Inst. Fourier (Grenoble) **75** (2025), no. 2, 769–828.
- [Huy23] Daniel Huybrechts, *The geometry of cubic hypersurfaces*, Cambridge Studies in Advanced Mathematics, vol. 206, Cambridge University Press, Cambridge, 2023.
- [Iva97] Nikolai V. Ivanov, *Automorphism of complexes of curves and of Teichmüller spaces*, Internat. Math. Res. Notices (1997), no. 14, 651–666.
- [Kan23] Jyoichi Kaneko, *The fundamental group of the complement of the Hesse configuration*, Internat. J. Math. **34** (2023), no. 4, Paper No. 2350014, 10.
- [Ker83] Steven P. Kerckhoff, *The Nielsen realization problem*, Ann. of Math. (2) **117** (1983), no. 2, 235–265.

- [KN63] Shoshichi Kobayashi and Katsumi Nomizu, *Foundations of differential geometry. Vol I*, Interscience Publishers (a division of John Wiley & Sons, Inc.), New York-London, 1963.
- [Kon00] Shigeyuki Kondō, *A complex hyperbolic structure for the moduli space of curves of genus three*, J. Reine Angew. Math. **525** (2000), 219–232.
- [Kon07] ———, *The moduli space of 5 points on $\mathbb{P}^1$ and K3 surfaces*, Arithmetic and geometry around hypergeometric functions, Progr. Math., vol. 260, Birkhäuser, Basel, 2007, pp. 189–206.
- [Kor99] Mustafa Korkmaz, *Automorphisms of complexes of curves on punctured spheres and on punctured tori*, Topology Appl. **95** (1999), no. 2, 85–111.
- [KR14] Ilya Kapovich and Kasra Rafi, *On hyperbolicity of free splitting and free factor complexes*, Groups Geom. Dyn. **8** (2014), no. 2, 391–414.
- [L17] Clara Löh, *Geometric group theory*, Universitext, Springer, 2017, An introduction.
- [Laz09] Radu Laza, *Deformations of singularities and variation of GIT quotients*, Trans. Amer. Math. Soc. **361** (2009), no. 4, 2109–2161.
- [Laz10] ———, *The moduli space of cubic fourfolds via the period map*, Ann. of Math. (2) **172** (2010), no. 1, 673–711.
- [Lee18] John M. Lee, *Introduction to Riemannian manifolds*, Graduate Texts in Mathematics, vol. 176, Springer, Cham, 2018, Second edition of [MR1468735].
- [Lib78] Anatoly Libgober, *On the fundamental group of the space of cubic surfaces*, Math. Z. **162** (1978), no. 1, 63–67.
- [LIP24] Claudio Llosa Isenrich and Pierre Py, *Subgroups of hyperbolic groups, finiteness properties and complex hyperbolic lattices*, Invent. Math. **235** (2024), no. 1, 233–254.
- [LLW22] Jun Li, Tian-Jun Li, and Weiwei Wu, *Symplectic Torelli groups of rational surfaces*, 2022, arXiv:2212.01873.
- [LO18] Radu Laza and Kieran G. O’Grady, *GIT versus Baily-Borel compactification for quartic K3 surfaces*, Geometry of moduli, Abel Symp., vol. 14, Springer, Cham, 2018, pp. 217–283. MR 3968048
- [LO19] Radu Laza and Kieran O’Grady, *Birational geometry of the moduli space of quartic K3 surfaces*, Compos. Math. **155** (2019), no. 9, 1655–1710.
- [Loo78] Eduard Looijenga, *On the semi-universal deformation of a simple-elliptic hypersurface singularity. II. The discriminant*, Topology **17** (1978), no. 1, 23–40.
- [Loo81] ———, *Rational surfaces with an anticanonical cycle*, Ann. of Math. (2) **114** (1981), no. 2, 267–322.
- [Loo84] E. Looijenga, *The smoothing components of a triangle singularity. II*, Math. Ann. **269** (1984), no. 3, 357–387.
- [Loo03a] Eduard Looijenga, *Compactifications defined by arrangements. I. The ball quotient case*, Duke Math. J. **118** (2003), no. 1, 151–187.
- [Loo03b] ———, *Compactifications defined by arrangements. II. Locally symmetric varieties of type IV*, Duke Math. J. **119** (2003), no. 3, 527–588.
- [Loo08] ———, *Artin groups and the fundamental groups of some moduli spaces*, J. Topol. **1** (2008), no. 1, 187–216.
- [Loo09] ———, *The period map for cubic fourfolds*, Invent. Math. **177** (2009), no. 1, 213–233.
- [Loo16] ———, *Moduli spaces and locally symmetric varieties*, Development of moduli theory—Kyoto 2013, Adv. Stud. Pure Math., vol. 69, Math. Soc. Japan, [Tokyo], 2016, pp. 33–75.
- [LS07] Eduard Looijenga and Rogier Swierstra, *The period map for cubic threefolds*, Compos. Math. **143** (2007), no. 4, 1037–1049.
- [Lüc10] Wolfgang Lück, *K- and L-theory of group rings*, Proceedings of the International Congress of Mathematicians. Volume II, Hindustan Book Agency, New Delhi, 2010, pp. 1071–1098.
- [LX19] Yuchen Liu and Chenyang Xu, *K-stability of cubic threefolds*, Duke Math. J. **168** (2019), no. 11, 2029–2073.
- [Man08] Jason Fox Manning, *Actions of certain arithmetic groups on Gromov hyperbolic spaces*, Algebr. Geom. Topol. **8** (2008), no. 3, 1371–1402.

- [Man25] Giorgio Mangioni, *Short hierarchically hyperbolic groups I: uncountably many coarse median structures*, 2025, arXiv:2410.09232.
- [Mas76] Howard Masur, *Extension of the Weil-Petersson metric to the boundary of Teichmüller space*, Duke Math. J. **43** (1976), no. 3, 623–635.
- [May77] Coy L. May, *Large automorphism groups of compact Klein surfaces with boundary. I*, Glasgow Math. J. **18** (1977), no. 1, 1–10.
- [Min18] Barry Minemeyer, *Real hyperbolic hyperplane complements in the complex hyperbolic plane*, Adv. Math. **338** (2018), 1038–1076.
- [MM99] Howard A. Masur and Yair N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149.
- [MM00] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. II. Hierarchical structure*, Geom. Funct. Anal. **10** (2000), no. 4, 902–974.
- [MMR90] J. Marsden, R. Montgomery, and T. Ratiu, *Reduction, symmetry, and phases in mechanics*, Mem. Amer. Math. Soc. **88** (1990), no. 436, iv+110.
- [MOT15] Shouhei Ma, Hisanori Ohashi, and Shingo Taki, *Rationality of the moduli spaces of Eisenstein K3 surfaces*, Trans. Amer. Math. Soc. **367** (2015), no. 12, 8643–8679.
- [MR08] Mahan Mj and Lawrence Reeves, *A combination theorem for strong relative hyperbolicity*, Geom. Topol. **12** (2008), no. 3, 1777–1798.
- [MS21] Joseph Maher and Saul Schleimer, *The compression body graph has infinite diameter*, Algebr. Geom. Topol. **21** (2021), no. 4, 1817–1856.
- [MSY92] Keiji Matsumoto, Takeshi Sasaki, and Masaaki Yoshida, *The monodromy of the period map of a 4-parameter family of K3 surfaces and the hypergeometric function of type (3, 6)*, Internat. J. Math. **3** (1992), no. 1, 1–164.
- [Nam85] Yukihiko Namikawa, *Periods of Enriques surfaces*, Math. Ann. **270** (1985), no. 2, 201–222.
- [Nar82] Isao Naruki, *Cross ratio variety as a moduli space of cubic surfaces*, Proc. London Math. Soc. (3) **45** (1982), no. 1, 1–30, With an appendix by Eduard Looijenga.
- [Nar90] ———, *The fundamental group of the complement for Klein’s arrangement of twenty-one lines*, Topology Appl. **34** (1990), no. 2, 167–181.
- [OSS16] Yuji Odaka, Cristiano Spotti, and Song Sun, *Compact moduli spaces of del Pezzo surfaces and Kähler-Einstein metrics*, J. Differential Geom. **102** (2016), no. 1, 127–172.
- [Par] John R. Parker, *Notes on Complex Hyperbolic Geometry*, <https://maths.dur.ac.uk/users/j.r.parker/img/NCHG.pdf>.
- [Pet26] Harry Petyt, *Mapping class groups are quasicubical*, Amer. J. Math. **148** (2026), no. 3, 771–806.
- [PP10] Jouni Parkkonen and Frédéric Paulin, *Prescribing the behaviour of geodesics in negative curvature*, Geom. Topol. **14** (2010), no. 1, 277–392.
- [PP16] D. Panov and A. Petrunin, *Ramification conjecture and Hirzebruch’s property of line arrangements*, Compos. Math. **152** (2016), no. 12, 2443–2460.
- [PS23] Harry Petyt and Davide Spriano, *Unbounded domains in hierarchically hyperbolic groups*, Groups Geom. Dyn. **17** (2023), no. 2, 479–500.
- [Pvv71] I. I. Pjateckiĭ-Šapiro and I. R. Šafarevič, *Torelli’s theorem for algebraic surfaces of type K3*, Izv. Akad. Nauk SSSR Ser. Mat. **35** (1971), 530–572.
- [Qui86] Frank Quinn, *Isotopy of 4-manifolds*, J. Differential Geom. **24** (1986), no. 3, 343–372.
- [RST23] Jacob Russell, Davide Spriano, and Hung Cong Tran, *Convexity in hierarchically hyperbolic spaces*, Algebr. Geom. Topol. **23** (2023), no. 3, 1167–1248.
- [Rus22] Jacob Russell, *From hierarchical to relative hyperbolicity*, Int. Math. Res. Not. IMRN (2022), no. 1, 575–624.
- [Sco78] Peter Scott, *Subgroups of surface groups are almost geometric*, J. London Math. Soc. (2) **17** (1978), no. 3, 555–565.
- [Sei08] Paul Seidel, *Lectures on four-dimensional Dehn twists*, Symplectic 4-manifolds and algebraic surfaces, Lecture Notes in Math., vol. 1938, Springer, Berlin, 2008, pp. 231–267.
- [Shi63] Goro Shimura, *On analytic families of polarized abelian varieties and automorphic functions*, Ann. of Math. (2) **78** (1963), 149–192.
- [Shi64] ———, *On purely transcendental fields of automorphic functions of several variable*, Osaka Math. J. **1** (1964), no. 1, 1–14.

- [Shi26] Ichiro Shimada, *Zariski pairs on cubic surfaces*, J. Knot Theory Ramifications **35** (2026), no. 4, Paper No. 2550083, 9.
- [Sis] Alessandro Sisto, *Lecture notes on Geometric Group Theory*, <https://alexsisto.net/LectureNotesGGT.pdf>.
- [Sis19] Alessandro Sisto, *What is a hierarchically hyperbolic space?*, Beyond hyperbolicity, London Math. Soc. Lecture Note Ser., vol. 454, Cambridge Univ. Press, Cambridge, 2019, pp. 117–148.
- [Sis25] Alessandro Sisto, *New tools in hierarchical hyperbolicity: A survey*, 2025, arXiv:2507.17546.
- [Smi15] Ivan Smith, *A symplectic prolegomenon*, Bull. Amer. Math. Soc. (N.S.) **52** (2015), no. 3, 415–464.
- [Smi22] Gleb Smirnov, *Symplectic mapping class groups of K3 surfaces and Seiberg-Witten invariants*, Geom. Funct. Anal. **32** (2022), no. 2, 280–301.
- [Smi23] ———, *Symplectic mapping class groups of blowups of tori*, J. Topol. **16** (2023), no. 3, 877–898.
- [Sor21] Ignat Soroko, *Artin groups of types F_4 and H_4 are not commensurable with that of type D_4* , Topology Appl. **300** (2021), Paper No. 107770, 15.
- [ST22] Matthew Stover and Domingo Toledo, *Residual finiteness for central extensions of lattices in $PU(n, 1)$ and negatively curved projective varieties*, Pure Appl. Math. Q. **18** (2022), no. 4, 1771–1797.
- [Tao25] Bingxue Tao, *An extension theorem for quasimorphisms*, 2025, arXiv:2511.21306.
- [Thu98] William P. Thurston, *Shapes of polyhedra and triangulations of the sphere*, The Epstein birthday schrift, Geom. Topol. Monogr., vol. 1, Geom. Topol. Publ., Coventry, 1998, pp. 511–549.
- [Tre16] Paula Tretkoff, *Complex ball quotients and line arrangements in the projective plane*, Mathematical Notes, vol. 51, Princeton University Press, Princeton, NJ, 2016, With an appendix by Hans-Christoph Im Hof.
- [Vö5] Jussi Väisälä, *Gromov hyperbolic spaces*, Expo. Math. **23** (2005), no. 3, 187–231.
- [Voi86] Claire Voisin, *Théorème de Torelli pour les cubiques de $\mathbf{P}^5$* , Invent. Math. **86** (1986), no. 3, 577–601.
- [Wri] Alex Wright, *Course notes for Math 797: Coarse geometry and Teichmüller theory*, <https://websites.umich.edu/~alexmw/Math797Notes.pdf>.
- [Wri26] ———, *What is... hierarchical hyperbolicity?*, 2026, arXiv:2601.15410.
- [Wu14] Weiwei Wu, *Exact Lagrangians in A_n -surface singularities*, Math. Ann. **359** (2014), no. 1-2, 153–168.
- [Yok02] Mutsumi Yokoyama, *Stability of cubic 3-folds*, Tokyo J. Math. **25** (2002), no. 1, 85–105.
- [Zha24] Junyan Zhao, *Compactifications of moduli of del Pezzo surfaces via line arrangement and K-stability*, Canad. J. Math. **76** (2024), no. 6, 2115–2135.