

EXTENSIONS OF VEECH GROUPS II: HIERARCHICAL HYPERBOLICITY AND QUASI-ISOMETRIC RIGIDITY

SPENCER DOWDALL, MATTHEW G. DURHAM, CHRISTOPHER J. LEININGER,
AND ALESSANDRO SISTO

ABSTRACT. We show that for any lattice Veech group in the mapping class group $\text{Mod}(S)$ of a closed surface S , the associated $\pi_1 S$ -extension group is a hierarchically hyperbolic group. As a consequence, we prove that any such extension group is quasi-isometrically rigid.

1. INTRODUCTION

This paper studies geometric properties of surface group extensions and how these relate to their defining subgroups of mapping class groups. Let S be a closed, connected, oriented surface of genus at least 2. Recall that a $\pi_1 S$ -extension of a group G is a short exact sequence of the form

$$1 \rightarrow \pi_1 S \rightarrow \Gamma \rightarrow G \rightarrow 1.$$

Such extensions are in bijective correspondence with monodromy homomorphisms from G to the extended mapping class group $\text{Mod}^\pm(S) \cong \text{Out}(\pi_1 S)$ of the surface. Alternatively, these groups Γ are precisely the fundamental groups of S -bundles.

Many advances in the study of mapping class groups have been motivated by a longstanding but incomplete analogy between hyperbolic space \mathbb{H}^n and the Teichmüller space $\mathcal{T}(S)$ of a surface. In the theory of Kleinian groups, a discrete group of isometries of \mathbb{H}^n is convex cocompact if it acts cocompactly on an invariant, convex subset. Farb and Mosher [FM02a] adapted this notion to mapping class groups by defining a subgroup $G \leq \text{Mod}^\pm(S)$ to be *convex cocompact* if it acts cocompactly on a quasi-convex subset of $\mathcal{T}(S)$. This has proven to be a fruitful concept with many interesting connections to, for example, the intrinsic geometry of the mapping class group [DT15, BBKL20], and its actions on the curve complex and the boundary of Teichmüller space [KL08a]. Most importantly, the work of Farb–Mosher [FM02a] and Hamenstädt [Ham] remarkably shows that an extension Γ as above is word hyperbolic if and only if the associated monodromy $G \rightarrow \text{Mod}^\pm(S)$ has finite kernel and convex cocompact image (see also [MS12]).

For Kleinian groups, convex cocompactness is a special case of a more prevalent phenomenon called *geometric finiteness*, which roughly amounts to acting cocompactly on a convex subset minus horoballs invariant by parabolic subgroups. In [Mos06], Mosher suggested this notion should have an analogous framework in mapping class groups that would extend the geometric connection with surface bundles to a larger class of examples. The prototypical candidates for geometric finiteness are the *lattice Veech subgroups*; these are special punctured-surface subgroups of $\text{Mod}(S)$ that arise naturally in the context of Teichmüller dynamics and whose corresponding S -bundles are amenable to study via techniques from flat geometry.

Our prequel paper [DDLS21] initiated an analysis of the $\pi_1 S$ -extensions associated to lattice Veech subgroups, with the main result being that each such extension Γ admits an action on a hyperbolic space \hat{E} that captures much of the geometry of Γ . Building on that work, the first main result of this paper is the following, which provides a concrete answer to [Mos06, Problem 6.2] for lattice Veech groups.

Theorem 1.1. *For any lattice Veech subgroup $G < \text{Mod}(S)$, the associated $\pi_1 S$ -extension group Γ of G is a hierarchically hyperbolic group.*

Hierarchical hyperbolicity means that in fact all the geometry of Γ is robustly encoded by hyperbolic spaces. This is exactly the sort of relaxed hyperbolicity for $\pi_1 S$ -extensions that one hopes should follow from a good definition of geometric finiteness in $\text{Mod}(S)$. Thus Theorem 1.1 suggests a possible general theory of geometric finiteness, which we expound upon in §1.4 below.

Hierarchical hyperbolicity has many strong consequences, some of which are detailed in §1.1 below. It also enables, via tools from [BHS21], the proof of our second main result, which answers [Mos06, Problem 5.4]:

Theorem 1.2. *For any lattice Veech group $G < \text{Mod}(S)$, the associated $\pi_1 S$ -extension group Γ of G is quasi-isometrically rigid.*

The rest of this introduction gives a more in-depth treatment of these results while elaborating on the concepts of, and connections between, hierarchical hyperbolicity, extensions of Veech groups, quasi-isometric rigidity, and geometric finiteness.

1.1. Hierarchical hyperbolicity. The notion of hierarchical hyperbolicity was defined by Behrstock, Hagen, and Sisto [BHS17b] and motivated by the seminal work of Masur and Minsky [MM00]. In short, it provides a framework and toolkit for understanding the coarse geometry of a space/group in terms of interrelated hyperbolic pieces. More precisely, a *hierarchically hyperbolic space (HHS)* structure on a metric space X is a collection of hyperbolic spaces $\{\mathcal{C}(W)\}_{W \in \mathfrak{S}}$, arranged in a hierarchical fashion, in which any pair are nested \sqsubseteq , orthogonal \perp , or transverse \pitchfork , along with Lipschitz projections to and between these spaces that together capture the coarse geometry of X . A *hierarchically hyperbolic group (HHG)* is then an HHS structure on a group that is equivariant with respect to an appropriate action on the union of hyperbolic spaces $\mathcal{C}(W)$. See §4 for details or [BHS17b, BHS19, Sis19] for many examples and further discussion.

Showing that a space/group is a hierarchically hyperbolic gives access to several results regarding, for example, a coarse median structure and quadratic isoperimetric function [Bow18, Bow13], asymptotic dimension [BHS17a], stable and quasiconvex subsets and subgroups [ABD21, RST18], quasiflats [BHS21], bordifications and automorphisms [DHS17], and quasi-isometric embeddings of nilpotent groups [BHS17b]. In particular, the following is an immediate consequence of Theorem 1.1.

Corollary 1.3. *Let $G < \text{Mod}(S)$ be any lattice Veech group and Γ the associated $\pi_1 S$ -extension group. Then:*

- (1) Γ has quadratic Dehn function [Bow13].
- (2) Γ is acylindrically hyperbolic and, moreover, its action on the \sqsubseteq -maximal hyperbolic space in the hierarchy is a universal acylindrical action [ABD21].

(3) Γ is semihyperbolic and thus has solvable conjugacy problem [DMS20, HHP20].

As discussed in §1.2 below, further information about Γ can be gleaned from the specific HHG structure constructed in proving Theorem 1.1. We note that the \sqsubseteq -maximal hyperbolic space of this structure, and thus the universal acylindrical action indicated in Corollary 1.3(2), is simply the space \hat{E} from [DDLS21].

1.2. The HHG structure on Γ . In order to describe the HHG structure more precisely and explain its connection to quasi-isometric rigidity in Theorem 1.2, we must first recall some of the structure of Veech groups and their extensions. Let $G < \text{Mod}(S)$ be a lattice Veech group and $\Gamma = \Gamma_G$ the associated extension group. First note that (up to finite index) Γ is naturally the fundamental group of an S -bundle \bar{E}/Γ over a compact surface with boundary (see §2 for details and notation). Each boundary component of \bar{E}/Γ is virtually the mapping torus of a multi-twist on S , and is thus a *graph manifold*: the tori in the JSJ decomposition are suspensions of the multi-twist curves.

Graph manifolds admit HHS structures [BHS19] where the maximal hyperbolic space is the Bass–Serre tree dual to the JSJ decomposition, and all other hyperbolic spaces are either quasi-lines or quasi-trees (obtained by coning off the boundaries of the universal covers of the base orbifolds of the Seifert pieces). The stabilizers of the vertices of the Bass–Serre trees are called *vertex subgroups*, and are precisely the fundamental groups of the Seifert pieces of the JSJ decomposition. We let \mathcal{V} denote the disjoint union of the vertices of all Bass–Serre trees associated to the boundary components of the universal cover \bar{E} of this S -bundle. Given $v, w \in \mathcal{V}$, we say that these vertices are *adjacent* if they are connected by an edge in the same Bass–Serre tree.

The HHG structure on the extension group Γ may now be described as follows:

Theorem 1.4. *Suppose $G < \text{Mod}(S)$ is a lattice Veech group with extension group Γ and let $\Upsilon_1, \dots, \Upsilon_k < \Gamma$ be representatives of the conjugacy classes of vertex subgroups. Then Γ admits an HHG structure with the following set of hyperbolic spaces and relations among them (ignoring those of diameter ≤ 4):*

- (1) *The maximal hyperbolic space \hat{E} is quasi-isometric to the Cayley graph of Γ coned off along the cosets of $\Upsilon_1, \dots, \Upsilon_k$ [DDLS21].*
- (2) *There is a quasi-tree v^{qt} and a quasi-line v^{ql} , for each $v \in \mathcal{V}$, and:*
 - (a) *For all $v \in \mathcal{V}$, $v^{qt} \perp v^{ql}$.*
 - (b) *For all $v, w \in \mathcal{V}$, if v and w are adjacent, then $w^{ql} \perp v^{ql}$ and $w^{ql} \sqsubseteq v^{qt}$.*
 - (c) *All other pairs are transverse.*

This description of the HHG structure readily leads to further consequences for Γ . For example, the maximal number of infinite-diameter pairwise orthogonal hyperbolic spaces is evidently 2. In view of [BHS17b, BHS21], we thus see that Γ is as “close to hyperbolic” as possible in that its quasi-flats are at worst 2-dimensional:

Corollary 1.5. *Each top-dimensional quasi-flat in Γ has dimension 2 and is contained in a finite-radius neighborhood of finitely many cosets of vertex subgroups.*

We note that quasiflats will be crucial for our proof of quasi-isometric rigidity, and we remark that the analogous statement for graph manifolds is due to Kapovich–Leeb [KL97].

Recall that an element of a group is a *generalized loxodromic* if it acts loxodromically under some acylindrical action on a hyperbolic space, and that a *universal*

acylindrical action on a hyperbolic space is one in which every generalized loxodromic acts loxodromically [ABD21]. It is shown in [Sis16] that a generalized loxodromic element g of a finitely generated group is necessarily *Morse*, meaning that in any finite-valence Cayley graph for the group, any (K, C) -quasi-geodesic with endpoints in the cyclic subgroup $\langle g \rangle$ stays within controlled distance $M = M(K, C)$ of $\langle g \rangle$. While being Morse is, in general, strictly weaker than being generalized loxodromic, these conditions are in fact equivalent in HHGs [ABD21, Theorem B].

In the case of our extension group Γ , it follows from Corollary 1.3(2) that the generalized loxodromics and Morse elements are precisely those elements acting loxodromically on \hat{E} . In [DDLS21, Theorem 1.1] we characterized these elements in terms of the vertex subgroups of Γ , thus yielding the following:

Corollary 1.6. *Let Γ be a lattice Veech group extension with vertex subgroups $\Upsilon_1, \dots, \Upsilon_k$ as in Theorem 1.4. The following are equivalent for an infinite order element $\gamma \in \Gamma$:*

- γ is not conjugate into any of the vertex subgroups Υ_i
- γ is a generalized loxodromic element of Γ
- γ is a Morse element of Γ .

1.3. Quasi-isometric rigidity. To state our rigidity theorem, first recall that Γ is (up to finite index) the fundamental group of an S -bundle \bar{E}/Γ over a compact surface with boundary. Here \bar{E} is a Γ -invariant truncation of the universal \hat{S} -bundle over the Teichmüller disk stabilized by the Veech group G . In particular, \bar{E} is quasi-isometric to Γ . Let $\text{Isom}(\bar{E})$ and $\text{QI}(\bar{E})$ denote the isometry and quasi-isometry groups of \bar{E} , respectively, and let $\text{Isom}_{\text{fib}}(\bar{E}) \leq \text{Isom}(\bar{E})$ denote the subgroup of isometries that map fibers to fibers.

Theorem 1.7. *There is an allowable truncation \bar{E} of E such that the natural homomorphisms $\text{Isom}_{\text{fib}}(\bar{E}) \rightarrow \text{Isom}(\bar{E}) \rightarrow \text{QI}(\bar{E}) \cong \text{QI}(\Gamma)$ are all isomorphisms, and $\Gamma \leq \text{Isom}(\bar{E}) \cong \text{QI}(\Gamma)$ has finite index.*

This is an analog, and indeed was motivated by, Farb and Mosher’s [FM02b] theorem that in the case of a surface group extension Γ_H associated to a Schottky subgroup H of $\text{Mod}(S)$, the natural homomorphism $\Gamma_H \rightarrow \text{QI}(\Gamma_H)$ is injective with finite cokernel. This rigidity also leads to the following strong *algebraic* consequence:

Corollary 1.8. *If H is any finitely generated group quasi-isometric to Γ , then H and Γ are weakly commensurable.*

In the statement, recall that two groups H_1, H_2 are *weakly commensurable* if there are finite normal subgroups $N_i \triangleleft H_i$ so that the quotients H_i/N_i have a pair of finite-index subgroups that are isomorphic to each other.

1.4. Motivation and Geometric Finiteness. Before outlining the paper and providing some ideas about the proofs, we provide some speculative discussion. For Kleinian groups—that is, discrete groups of isometries of hyperbolic 3-space—the notion of geometric finiteness is important in the deformation theory of hyperbolic 3-manifolds by the work of Ahlfors [Ahl66] and Greenberg [Gre66]. While the definition has many formulations (see [Mar74, Mas70, Thu86, Bow93]), roughly speaking a group is geometrically finite if it acts cocompactly on a convex subset of hyperbolic 3-space minus a collection of horoballs that are invariant by parabolic subgroups. When there are no parabolic subgroups, geometric finiteness reduces

to convex cocompactness: a cocompact action on a convex subset of hyperbolic 3-space.

While there is no deformation theory for subgroups of mapping class groups, Farb and Mosher [FM02a] introduced a notion of convex cocompactness for $G < \text{Mod}(S)$ in terms of the action on Teichmüller space $\mathcal{T}(S)$. Their definition requires that G acts cocompactly on a quasi-convex subset for the Teichmüller metric, while Kent and Leininger later proved a variety of equivalent formulations analogous to the Kleinian setting [KL07, KL08a, KL08b]. Farb and Mosher proved that convex cocompactness is equivalent to hyperbolicity of the associated extension group Γ_G (with monodromy given by inclusion) when G is virtually free. This equivalence was later proven in general by Hamenstädt [Ham] (see also Mj–Sardar [MS12]), though at the moment the only known examples are virtually free.

The coarse nature of Farb and Mosher’s formulation reflects the fact that the Teichmüller metric is far less well-behaved than that of hyperbolic 3-space. Quasi-convexity in the definition is meant to help with the lack of nice *local* behavior of the Teichmüller metric. It also helps with the *global* lack of Gromov hyperbolicity (see Masur–Wolf [MW95]), as cocompactness of the action ensures that the quasi-convex subset in the definition is Gromov hyperbolic (see Kent–Leininger [KL08a], Minsky [Min96b], and Rafi [Raf14]).

The inclusion of reducible/parabolic mapping classes in a subgroup $G < \text{Mod}(S)$ brings the *thin parts* of $\mathcal{T}(S)$ into consideration; these subspaces contain higher rank quasi-flats and even exhibit aspects of positive curvature (see Minsky [Min96a]). This is a main reason why extending the notion of convex cocompactness to geometric finiteness is complicated. These complications are somewhat mitigated in the case of lattice Veech groups. Such subgroups are stabilizers of isometrically and totally geodesically embedded hyperbolic planes, called *Teichmüller disks*, that have finite area quotients. Thus, the intrinsic hyperbolic geometry agrees with the extrinsic Teichmüller geometry, and as a group of isometries of the hyperbolic plane, a lattice Veech group *is* geometrically finite. This is why these subgroups serve as a test case for geometric finiteness in the mapping class group. This is also why a *subgroup* of a Veech group is convex cocompact in $\text{Mod}(S)$ if and only if it is convex cocompact as a group of isometries of the hyperbolic plane (which also happens if and only if it is finitely generated and contains no parabolic elements).

The action of $\text{Mod}(S)$ on the *curve graph*, which is Gromov hyperbolic by work of Masur–Minsky [MM99], provides an additional model for these considerations. Specifically, convex cocompactness is equivalent to the orbit map to the curve graph $\mathcal{C}(S)$ being a quasi-isometric embedding with respect to the word metric from a finite generating set (see Kent–Leininger [KL08a] and Hamenstädt [Ham]). Viewing geometric finiteness as a kind of “relative convex cocompactness” for Kleinian groups suggests an interesting connection with the curve complex formulation. The connection is best illustrated by the following theorem of Tang [Tan19].

Theorem 1.9 (Tang). *For any lattice Veech group $G < \text{Mod}(S)$ stabilizing a Teichmüller disk $D \subset \mathcal{T}(S)$, there is a G -equivariant quasi-isometric embedding $D^{el} \rightarrow \mathcal{C}(S)$, where D^{el} is the path metric space obtained from D by coning off the G -invariant family of horoballs in which D ventures into the thin parts of $\mathcal{T}(S)$.*

Farb [Far98] showed that non-cocompact lattices in the group of isometries of hyperbolic space are *relatively hyperbolic* relative to the parabolic subgroups. For Veech groups, the space D^{el} is quasi-isometric to the (hyperbolic) coned off Cayley

graph, illustrating (part of) the relative hyperbolicity of G . We thus propose a kind of “qualified” notion of geometric finiteness with this in mind:

Definition 1.10 (Parabolic geometric finiteness). A finitely generated subgroup $G < \text{Mod}(S)$ is *parabolically geometrically finite* if G is relatively hyperbolic, relative to a (possibly trivial) collection of subgroups $\mathcal{H} = \{H_1, \dots, H_k\}$, and

- (1) H_i contains a finite index, abelian subgroup consisting entirely of multi-twists, for each $1 \leq i \leq k$; and
- (2) the coned off Cayley graph G -equivariantly and quasi-isometrically embeds into $\mathcal{C}(S)$.

When $\mathcal{H} = \{\{id\}\}$, we note that the condition is equivalent to G being convex cocompact. By Theorem 1.9, lattice Veech groups are parabolically geometrically finite. In fact, Tang’s result is more general and implies that any finitely generated Veech group satisfies this definition. These examples are all virtually free, but other examples include the combination subgroups of Leininger–Reid [LR06], which are isomorphic to fundamental groups of closed surfaces of higher genus, and free products of higher rank abelian groups constructed by Loa [Loa21].

In view of Theorem 1.1, one might formulate the following.

Conjecture 1.11. *Let $G < \text{Mod}(S)$ be parabolically geometrically finite. Then the $\pi_1 S$ -extension group Γ of G is a hierarchically hyperbolic group.*

We view Definition 1.10 as only a qualified formulation because there are many subgroups of $\text{Mod}(S)$ that are not relatively hyperbolic but are nevertheless candidates for being geometrically finite in some sense. It is possible that there are different types of geometric finiteness for subgroups of mapping class groups, with Definition 1.10 being among the most restrictive. Other notions might include an HHS structure on the subgroup which is compatible with the ambient one on $\text{Mod}(S)$ (e.g., hierarchical quasiconvexity [BHS19]). From this perspective, some candidate subgroups that may be considered geometrically finite include:

- the whole group $\text{Mod}(S)$;
- multi-curve stabilizers;
- the right-angled Artin subgroups of mapping class groups constructed in [CLM12, Kob12, Run20];
- free and amalgamated products of other examples.

Question 1.12. For each example group $G \leq \text{Mod}(S)$ above, is the associated extension Γ_G a hierarchically hyperbolic group?

We note that the answer is ‘yes’ for the first example, since the extension group is the mapping class group of the surface S with a puncture. Moreover, since our work on this subject first appeared, Russell [Rus21] addressed the second example by proving extensions of multicurve stabilizers are hierarchically hyperbolic groups.

1.5. Outline and proofs. Let us briefly outline the paper and comment on the main structure of the proofs. In §2 we review necessary background material and introduce the objects and notation that will be used throughout the paper. In particular, we define the spaces E and \bar{E} , the latter being a quasi-isometric model for the Veech group extension Γ , as well as the hyperbolic collapsed space \hat{E} . All of these were constructed in [DDLS21].

In §§3–4 we prove that the extension group Γ is hierarchically hyperbolic by utilizing a combinatorial criterion from [BHMS20]. Besides hyperbolicity of \hat{E} , the other hard part of the criterion is an analogue of Bowditch’s fineness condition from the context of relative hyperbolicity. Its geometric interpretation is roughly that two cosets of vertex subgroups as above have bounded coarse intersection, aside from the “obvious” exception when the cosets correspond to vertices of the same Bass–Serre tree within distance 2 of each other. To this end, in §3 we associate to each vertex $v \in \mathcal{V}$ a spine bundle $\Theta^v \subset \bar{E}$, which corresponds to a Seifert piece of the JSJ decomposition of the peripheral graph manifold, along with a pair of hyperbolic spaces \mathcal{K}^v and Ξ^v that will figure into the HHS structure on Γ . The space \mathcal{K}^v is obtained via a quasimorphism constructed using the Seifert fibered structure following ideas in forthcoming work of the fourth author with Hagen, Russell, and Spriano [HRSS21], while Ξ^v is coarsely obtained by coning off boundary components of the universal covers of the base 2–orbifold of this Seifert fibered manifold. We then appeal to the flat geometry of the fibers of E to construct and study certain projection maps

$$\begin{array}{ccccc}
 & & \bar{E} & & \\
 & \swarrow & & \searrow & \\
 & \Lambda^v & & \xi^v & \\
 & \swarrow & \downarrow \Pi^v & \searrow & \\
 \mathcal{K}^v & \longleftarrow \lambda^v & \Theta^v & \longrightarrow i^v & \Xi^v
 \end{array}$$

and prove that various pairs of subspaces of \bar{E} have bounded projection onto each other (Proposition 3.19).

In §4, we begin assembling the combinatorial objects necessary to apply the HHG criterion from [BHMS20], which involves both combinatorial and geometric aspects. The first step involves the construction of a natural flag complex \mathcal{X} containing the union of the Bass–Serre trees, together with appropriate “subjoins” with the union of all \mathcal{K}^v , over $v \in \mathcal{V}$. Next, we use the geometry of \bar{E} to construct a certain graph \mathcal{W} whose vertices are maximal simplices of \mathcal{X} and on which Γ acts metrically properly and coboundedly. The remainder of this section is devoted to verifying the necessary combinatorial conditions as well as translating the facts about \mathcal{K}^v and Ξ^v and the projections described above into proofs of the necessary geometric conditions. We note that in the combinatorial HHG setup, the complex \mathcal{X} comes with its own hierarchy projections between the induced hyperbolic spaces (Definitions 4.9–4.10), which may be different than the projections to \mathcal{K}^v and Ξ^v .

In §5 we prove our QI-rigidity result Theorem 1.7. The starting point is the hierarchical hyperbolicity of Γ provided by Theorem 1.4, as it gives access to the results and arguments in [BHS21] about the preservation of quasi-isometrically embedded flats. Every collection of pairwise orthogonal hyperbolic spaces in an HHG determines a natural product subspace, with the maximal *standard* quasi-isometrically embedded flats (or orthants) arising inside such subspaces as products of quasi-lines in a maximal collection of pairwise orthogonal hyperbolic spaces of the HHG. Theorem A of [BHS21] states that a quasi-isometry of an HHS preserves the structure of its quasi-flats and takes any maximal quasi-flat within bounded Hausdorff distance of the union of standard maximal orthants. The maximal quasi-flats in the HHG structure on \bar{E} , namely the 2–dimensional flats indicated in Corollary 1.5, are encoded by certain *strip bundles* that, roughly, correspond to flats in the peripheral graph manifolds. We use the preservation of the maximal quasi-flats to derive

coarse preservation of these strip bundles, which we then upgrade to coarse preservation of the fibers (§5.1). By using tools of flat geometry from [BL18, DELS18], we then show any quasi-isometry induces an *affine* homeomorphism of any fiber to itself (§§5.2–5.3) and moreover that this assignment is injective (§5.4). Finally, we show this association is an isomorphism by proving (§5.5) that every affine homeomorphism of a fiber induces an isometry and hence quasi-isometry of \bar{E} . Quasi-isometric rigidity and its algebraic consequence Corollary 1.8 are then easily obtained in §5.6.

Acknowledgments. The authors would like to thank MSRI and its Fall 2016 program on *Geometric Group Theory*, where this work began. We also gratefully acknowledge NSF grants DMS 1107452, 1107263, 1107367 (the GEAR Network) for supporting travel related to this project. Dowdall was partially supported by NSF grants DMS-1711089 and DMS-2005368. Durham was partially supported by NSF grant DMS-1906487. Leininger was partially supported by NSF grants DMS-1510034, DMS-1811518, and DMS-2106419. Sisto was partially supported by the Swiss National Science Foundation (grant #182186).

2. SETUP: THE GROUPS AND SPACES

Here we briefly recall the basic set up from [DDLS21] which we will use throughout the remainder of the paper. We refer the reader to Sections 2 and 3 of that paper for details and precise references.

2.1. Flat metrics and Veech groups. Fix a closed surface of genus at least 2, a complex structures X_0 (viewed as a point in the Teichmüller space $\mathcal{T}(S)$), and a nonzero holomorphic quadratic differential q on (S, X_0) . Integrating a square root of q determines preferred coordinates on (S, X_0) for q which defines a translation structure (in the complement of the isolated zeros of q). We also write q for the associated flat metric defined by the half-translation structure (though the metric only determines the half-translation structure or quadratic differential up to a complex scalar multiple). This metric is a non-positively curved Euclidean cone metric, with cone singularities at the zeros of q . The orbit of (X_0, q) under the natural $\mathrm{SL}_2(\mathbb{R})$ action on quadratic differentials projects to a *Teichmüller disk*, $D = D_q \subset \mathcal{T}(S)$, which we equip with its Poincaré metric ρ . The circle at infinity of D is naturally identified with the projective space of directions, $\mathbb{P}^1(q)$, in the tangent space of any nonsingular point of q . For $\alpha \in \mathbb{P}^1(q)$, we write $\mathcal{F}(\alpha)$ for the singular foliation by geodesics in direction α .

We assume that the associated Veech group $G = G_q$ is a lattice—recall that G can be viewed as the stabilizer in the mapping class group of S of D as well as the affine group of q , and the lattice assumption is equivalent to requiring the quotient orbifold D/G to have finite ρ -area. The parabolic fixed points in the circle at infinity form a subset we denote $\mathcal{P} \subset \mathbb{P}^1(q)$. This subset corresponds precisely to the completely periodic directions for the flat metric q ; that is, the directions α for which the foliation $\mathcal{F}(\alpha)$ decomposes S into cylinders foliated by q -geodesic core circles. The boundaries of these cylinders are q -saddle connections (q -geodesic segments connecting pairs of cone points, with no cone points in their interior), and by the Veech Dichotomy, every saddle connection is in a direction in \mathcal{P} . We

let $\{B_\alpha\}_{\alpha \in \mathcal{P}}$ denote any G -invariant, 1-separated set of horoballs in D and let

$$\bar{D} = D \setminus \bigcup_{\alpha \in \mathcal{P}} \text{int}(B_\alpha)$$

be the G -invariant subspace obtained by removing these horoballs. We write $\bar{\rho}$ for the induced path metric on \bar{D} . Finally, we let

$$p: D \rightarrow \hat{D}$$

be the G -equivariant quotient obtained by collapsing each horoball B_α to a point, for $\alpha \in \mathcal{P}$. There is a natural path metric $\hat{\rho}$ on \hat{D} so that p is 1-Lipschitz and is a local isometry at every point not in one of the horoballs.

We will also make use of the closest point projection to the horoball

$$c_\alpha: D \rightarrow B_\alpha$$

for each $\alpha \in \mathcal{P}$.

2.2. The bundles E and \bar{E} . For each point $X \in D$, we let q_X denote the associated flat metric or quadratic differential (defined up to scalar multiplication) on S . The space of interest E is a bundle over D ,

$$\pi: E \rightarrow D,$$

for which the fiber E_X over $X \in D$ is naturally identified with the universal cover \tilde{S} of S , equipped with the pull-back complex structure X and quadratic differential/flat metric q_X . We write $\mathcal{B}_\alpha = \pi^{-1}(B_\alpha)$ for $\alpha \in \mathcal{P}$.

For any $X, Y \in D$, the Teichmüller map between these complex structures has initial and terminal quadratic differentials q_X and q_Y (up to scalar multiple) and this map lifts to a canonical affine map between the fibers $f_{Y,X}: E_X \rightarrow E_Y$. These maps satisfy $f_{Z,X} = f_{Z,Y} f_{Y,X}$ for all $X, Y, Z \in D$, and for any $X \in D$, assemble to a map $f_X: E \rightarrow E_X$ defined by $f_X(y) = f_{X,\pi(y)}(y)$. Moreover, for any $X, Y \in D$, $f_{Y,X}$ is $e^{\rho(X,Y)}$ -bi-Lipschitz. We use the maps f_{X,X_0} to identify $\mathbb{P}^1(q) \cong \mathbb{P}^1(q_X)$ for all $X \in D$.

The fiber over X_0 is denoted $E_0 = E_{X_0}$ and the maps $f_0 = f_{X_0}: E \rightarrow E_0$ and $\pi: E \rightarrow D$ are projections on the factors in a product structure $E \cong D \times E_0 \cong D \times \tilde{S}$. For $x \in E$, we write $D_x = f_{\pi(x)}^{-1}(x)$, which is just the slice $D \times \{f_0(x)\}$ in the product structure. The affine maps $f_{Y,X}$ sends the cone points Σ_X of E_X to the cone points Σ_Y of E_Y . Consequently, the union of all singular points

$$\Sigma = \bigcup_{X \in D} \Sigma_X$$

is a locally finite union of disks D_x , one for each $x \in \Sigma_0 = \Sigma_{X_0}$.

We give the space E a singular Riemannian metric d which is the flat metric on each fiber E_X and the Poincaré metric on each disk D_x so that at each smooth point of intersection, the tangent planes are orthogonal. The singular locus of this metric is precisely Σ . Each disk D_x is isometrically embedded since π is a 1-Lipschitz map, and hence restricts to an isometry $\pi|_{D_x}: D_x \rightarrow D$. The metric on $E \setminus \Sigma$ is in fact a locally homogeneous metric, modeled on a four-dimensional, Thurston-type geometry; see [DDLS21, §5].

The extension group Γ acts on E by bundle maps with the kernel $\pi_1 S < \Gamma$ of the projection to G acting trivially on D and by covering transformation on each fiber E_X . We set $\bar{E} = \pi^{-1}(\bar{D}) \subset E$, and write $\bar{\pi}: \bar{E} \rightarrow \bar{D}$. When convenient to do

so, we put “bars” over objects associated to \bar{D} or \bar{E} , e.g. $\bar{D}_x = D_x \cap \bar{E}$, $\bar{p}: \bar{D} \rightarrow \hat{D}$, etc. In particular, we write \bar{d} for the induced path metric on $\bar{E} \subset E$, induced from the metric on E described above.

For any $\alpha \in \mathcal{P}$, the closest point projection $c_\alpha: D \rightarrow B_\alpha$ has a useful “lift” $f_\alpha: E \rightarrow \mathcal{B}_\alpha$, defined by

$$f_\alpha(x) = f_{c_\alpha(\pi(x))}(x),$$

for any $x \in \bar{E}$. That is, f_α maps each fiber E_X via the map $f_{Y,X}$ to E_Y , where $Y = c_\alpha(X)$ is the image of the closest point projection to B_α of X in D .

2.3. The hyperbolic space \hat{E} . The quotient $p: D \rightarrow \hat{D}$ is the descent of a quotient $P: E \rightarrow \hat{E}$ which we now describe. First, for each $\alpha \in \mathcal{P}$, the foliation $\mathcal{F}(\alpha)$ lifts to a foliation on E_0 in direction α , and hence on any fiber E_X by push-forward via the map f_{X,X_0} , also in direction α (via the identification $\mathbb{P}^1(q) \cong \mathbb{P}^1(q_X)$). There is a natural transverse measure coming from the flat metric on X . Given $\alpha \in \mathcal{P}$, we fix some $X_\alpha \in \partial B_\alpha$ and let T_α be the dual simplicial \mathbb{R} -tree to this measured foliation in direction α on E_{X_α} , and we let $t_\alpha: E \rightarrow T_\alpha$ be the composition of the leafspace projection $E_{X_\alpha} \rightarrow T_\alpha$ with the map $f_{X_\alpha}: E \rightarrow E_{X_\alpha}$.

Now we define $P: E \rightarrow \hat{E}$ to be the quotient space obtained by collapsing the subset \mathcal{B}_α to T_α via $t_\alpha|_{\mathcal{B}_\alpha}$ for each $\alpha \in \mathcal{P}$. We also write $\bar{P} = P|_{\bar{E}}: \bar{E} \rightarrow \hat{E}$. The maps P and \bar{P} descend to the maps p and \bar{p} , and the map π determines maps $\hat{\pi}$ and $\bar{\pi}$, which all fit into the following commutative diagram.

$$\begin{array}{ccccc} & & E & & \\ & \nearrow & \downarrow \pi & \searrow P & \\ \bar{E} & \xrightarrow{\bar{P}} & D & \xrightarrow{p} & \hat{E} \\ & \searrow \bar{\pi} & & \nearrow \hat{\pi} & \\ \bar{D} & \xrightarrow{\bar{p}} & \hat{D} & & \end{array}$$

A metric \hat{d} on \hat{E} is determined by \bar{d} on \bar{E} and the map \bar{P} . The main facts about this metric are summarized in the following theorem; see [DDLS21, Theorem 1.1, Lemma 3.2].

Theorem 2.1. *There is a Gromov hyperbolic path metric \hat{d} on \hat{E} so that $\bar{P}: \bar{E} \rightarrow \hat{E}$ is 1-Lipschitz and is a local isometry at every point $x \in \bar{E} - \partial \bar{E}$. Furthermore, for every $\alpha \in \mathcal{P}$,*

- *The induced path metric on $P(\partial \mathcal{B}_\alpha) = T_\alpha$ is the \mathbb{R} -tree metric determined by the transverse measure on the foliation of E_{X_α} in direction α .*
- *The subspace topology on $T_\alpha \subset \hat{E}$ agrees with the \mathbb{R} -tree topology on T_α .*

Remark 2.2. The underlying simplicial tree T_α is precisely the Bass-Serre tree dual to the splitting of $\pi_1 S$ defined by the cores of the cylinders of $\mathcal{F}(\alpha)$ on S .

For each $x \in E$, we denote the image of D_x in \hat{E} by \hat{D}_x , which is obtained by collapsing $\mathcal{B}_\alpha \cap D_x$ to a point, for each $\alpha \in \mathcal{P}$. Consequently, $\hat{\pi}|_{\hat{D}_x}: \hat{D}_x \rightarrow \hat{D}$ is a bijection, and so each \hat{D}_x , with its path metric, is isometric to \hat{D} and isometrically embedded in \hat{E} . We call objects in E , \bar{E} , and \hat{E} *vertical* if they are contained in a fiber of π , $\bar{\pi}$, or $\hat{\pi}$, respectively, and *horizontal* if they are contained in D_x , \bar{D}_x , or \hat{D}_x , for some $x \in E$, \bar{E} .

2.4. Vertices, spines, and spine bundles. We will write $\mathcal{V} \subset \hat{E}$ for the union over all $\alpha \in \mathcal{P}$ of all vertices of T_α . We will simultaneously view \mathcal{V} as both a subset of \hat{E} and abstractly as an indexing set that will be used in sections §3–4 to develop an HHS structure on \bar{E} . Since each vertex belongs to a unique tree, and since the trees are indexed by $\alpha \in \mathcal{P}$, we obtain a map $\alpha: \mathcal{V} \rightarrow \mathcal{P}$ so that v is a vertex of $T_{\alpha(v)}$. For convenience, we also write $B_v = B_{\alpha(v)}$, $\partial B_v = \partial B_{\alpha(v)}$, etc for each $v \in \mathcal{V}$, and write $c_v = c_{\alpha(v)}$ for the ρ -closest point projection $D \rightarrow B_v$.

For $v, w \in \mathcal{V}$, we write $v \parallel w$ if $\alpha(v) = \alpha(w)$. Then define $d_{\text{tree}}(v, w) \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ to be the combinatorial (integer valued) distance in the simplicial tree $T_{\alpha(v)} = T_{\alpha(w)}$ when $v \parallel w$ (as opposed to the distance from the \mathbb{R} -tree metric) and to equal ∞ when $v \not\parallel w$.

Given $\alpha \in \mathcal{P}$, $X \in D$, and $v \in T_\alpha$, the v -*spine* in E_X is the subspace

$$\theta_X^v = (P \circ f_{X_\alpha, X})^{-1}(v) = t_\alpha^{-1}(v) \cap E_X.$$

The v -spine θ_X^v is the union of the saddle connections on the fiber E_X in direction α that project to v by t_α . When $d_{\text{tree}}(v, w) = 1$ (and hence v, w are adjacent in the same tree T_α) there is a unique component of $E_X \setminus (\theta_X^v \cup \theta_X^w)$ whose closure is an infinite strip, $\mathbb{R} \times [a, b]$, that covers a maximal cylinder in the quotient $E_X/\pi_1 S = (S, X, q_X)$ in the direction α . We let Θ_X^v be the union of θ_X^v and all such strips defined by $w \in T_\alpha$ with $d_{\text{tree}}(v, w) = 1$. We call Θ_X^v the *thickened v -spine in E_X* . Observe that the affine map $f_{Y, X}$ maps θ_X^v and Θ_X^v to θ_Y^v and Θ_Y^v , respectively, for all $X, Y \in D$. Finally, we write

$$\theta^v = \bigcup_{X \in \partial B_v} \theta_X^v \quad \Theta^v = \bigcup_{X \in \partial B_v} \Theta_X^v.$$

These spaces are bundles over ∂B_v which we call, respectively, the v -*spine bundle* and the *thickened v -spine bundle*.

2.5. Schematic of the space \bar{E} and its important pieces. Figure 1 is a cartoon of the bundle \bar{E} over the truncated Teichmüller disk \bar{D} . We have tried to highlight some of the key features of \bar{E} which are relevant to this paper.

- (a) The stabilizer of a horoball based at a point $\beta \in \mathcal{P}$ is virtually cyclic, generated by a multitwist τ_α acting as a parabolic on D . The base point X_β on the horocycle based at β and its image are shown.
- (b) The bundle over the boundary horocycle based at β is shown. This is the universal cover, $\partial \mathcal{B}_\beta$, of a graph manifold which is the mapping torus of τ_β . Two fibers E_{X_β} and $E_{\tau_\beta(X_\beta)}$ are shown with the effect on a part of a spine (in green) in some other direction illustrating the sheering in strips after applying τ_β .
- (c) This is another horoball in some direction α , with the chosen basepoint X_α and its horocycle ∂B_α .
- (d) The spine $\theta_{X_\alpha}^v$ in direction α is shown in red, corresponding to a vertex $v \in T_\alpha$. The thickened spine $\Theta_{X_\alpha}^v$ is indicated in lavender. Spines for vertices of T_α adjacent to v meet $\Theta_{X_\alpha}^v$ along lines in $\partial \Theta_{X_\alpha}^v$ and are shown in various other colors.
- (e) The restriction of $t_\alpha: E \rightarrow T_\alpha$ to E_{X_α} collapses each spine $\theta_{X_\alpha}^w$ or strip in direction α to the corresponding vertex w or edge the Bass-Serre tree T_α . The space \hat{E} is formed by collapsing \mathcal{B}_α to T_α via t_α .




FIGURE 1. A schematic of \bar{E} and various key features of it.

2.6. Some technical lemmas and coarse geometry. Here we briefly recall some basic facts about the setup above proved in [DDLS21] as well as some useful coarse geometric facts. The first fact is the following; see [DDLS21, Lemma 3.4].

Lemma 2.3. *There exists a constant $M > 0$ such that for each $v \in \mathcal{V}$ and $X \in \partial B_v$, every saddle connection in θ_X^v has length at most M and every strip in Θ_X^v has width at most M . In particular, for points $X \in \partial B_\alpha$, the saddle connections and strips of E_X in direction $\alpha \in \mathcal{P}$ have, respectively, uniformly bounded lengths and widths.*

Every connected graph can be made into a geodesic metric space by locally isometrically identifying each edge with a unit interval. We will need the following well-known result (for a proof of this version, see [DDLS21, Proposition 2.1]).

Proposition 2.4. *Let Ω be a path metric space and $\Upsilon \subset \Omega$ an R -dense subset for some $R > 0$. For any $R' > 3R$, consider a graph \mathcal{G} with vertex set Υ such that:*

- *all pairs of elements of Υ within distance $3R$ are joined by an edge in \mathcal{G} ,*
- *if an edge in \mathcal{G} joins points $w, w' \in \Upsilon$, then $d_\Omega(w, w') \leq R'$.*

Then the inclusion of Υ into Ω extends to a quasi-isometry $\mathcal{G} \rightarrow \Omega$.

The following criterion for a graph to be a quasi-tree is well-known, and an easy consequence of Manning's bottleneck criterion [Man05]. We include a proof for completeness.

Proposition 2.5. *Let X be a graph, and suppose that there exists a constant B with the following property: For each pair of vertices w, w' there exists an edge path*

$\gamma(w, w')$ from w to w' so that for any vertex v on $\gamma(w, w')$, any path from w to w' intersects the ball of radius B around v . Then X is quasi-isometric to a tree, with quasi-isometry constants depending on B only.

Proof. We check that [Man05, Theorem 4.6] applies; that is, we check the following property. For any two vertices $w, w' \in X$, there is a midpoint $m(w, w')$ between w and w' so that any path from w to w' passes within distance $B' = B'(B)$ of $m(w, w')$. (The uniformity in the quasi-isometry comes from the proof of Manning's theorem, see [Man05, page 1170].)

Consider any geodesic α from w to w' , and let $m = m(w, w')$ be its midpoint. We will show that m lies within distance $2B + 1$ of a vertex of $\gamma = \gamma(w, w')$, so that we can take $B' = 3B + 1$.

Indeed, suppose by contradiction that this is not the case. Let $w = w_0, \dots, w_n = w'$ be the vertices of γ (in the order in which they appear along γ), and let $d_i = d(w, w_i)$, so that $|d_{i+1} - d_i| \leq 1$. Each w_i lies within distance B of some point p_i on α which must satisfy $d(p_i, m) \geq B + 1$. In particular, we have that every d_i satisfies either $d_i \leq d(w, w')/2 - 1$ or $d_i \geq d(w, w')/2 + 1$. Since $d_0 = 0$ and $d_n = d(w, w')$, we cannot have $|d_{i+1} - d_i| \leq 1$ for all $0 \leq i \leq n - 1$, a contradiction. \square

We end with a few definitions from coarse geometry which may not be completely standard but will appear in the next two sections. Given two metrics d and d' on a set X , we say that d is *coarsely bounded* by d' if there exists a monotone function $N: [0, \infty) \rightarrow [0, \infty)$ so that $d(x, y) \leq N(d'(x, y))$, for all $x, y \in X$. If d is coarsely bounded by d' and d' is coarsely bounded by d , we say that d and d' are *coarsely equivalent*. An isometric action of a group H on a metric space Y is *metrically proper* if for any $R > 0$ and any point $y \in Y$, there are at most finitely many elements $h \in H$ for which $h \cdot B(y, R) \cap B(y, R) \neq \emptyset$. For proper geodesic spaces, this is equivalent to acting properly discontinuously. If there exists y, R so that $H \cdot B(y, R) = Y$, then we say that the action is cobounded, and for proper geodesic metric spaces this is equivalent to acting cocompactly.

3. PROJECTIONS AND VERTEX SPACES

An HHS structure on a metric space consists of certain additional data, most importantly a collection of hyperbolic spaces together with projection maps to each space. For the HHS structure that we will build on (Cayley graphs of) Γ , the hyperbolic spaces will (up to quasi-isometry) be the space \hat{E} from [DDLS21] (see §2.3) and the spaces \mathcal{K}^v and Ξ^v introduced in this section, where v varies over all vertices of the trees T_α . Morally, the projections will be given by the maps Λ^v and ξ^v that we study below. However, to prove hierarchical hyperbolicity we will use a criterion from [BHMS20] which does not require actually defining projections, but nevertheless provides them. Still, the maps Λ^v and ξ^v will play a crucial role in proving this criterion applies.

We will establish properties of Λ^v and ξ^v that are reminiscent of subsurface projections or of closest-point projections to peripheral sets in relatively hyperbolic spaces/groups; these are summarized in Proposition 3.19. Essentially, these same properties would be needed if we wanted to construct an HHS structure on Γ directly without using [BHMS20].

From a technical point of view, we would like to draw attention to Lemma 3.13, which is the crucial lemma that ensures that the projections behave as desired and

that various subspaces have bounded projections. Roughly, the lemma says that closest-point projections to a spine do not vary much under affine deformations.

In what follows, we will write d_{Θ^v} and $d_{\partial\mathcal{B}_\alpha}$ for the path metrics on Θ^v and $\partial\mathcal{B}_\alpha$ induced from \bar{d} . Using the map $f_\alpha: \bar{E} \rightarrow \partial\mathcal{B}_\alpha$, it is straightforward to see that $d_{\partial\mathcal{B}_\alpha}$ is uniformly coarsely equivalent to the subspace metric from \bar{d} : in fact, $\bar{d} \leq d_{\partial\mathcal{B}_\alpha} \leq e^{\bar{d}}$. The same is true for d_{Θ^v} , which follows from the fact that the inclusion of Θ^v into $\partial\mathcal{B}_\alpha$ is a quasi-isometric embedding with respect to the path metrics (see below).

Associated to each $v \in \mathcal{V}$ we will be considering two types of projections. These projections have a single projection $\Pi^v: \bar{\Sigma} \rightarrow \Theta^v$ as a common ingredient. It is convenient to analyze Π^v via an auxiliary map which serves as a kind of fiberwise closest point projection that survives affine deformations, and which we call the *window map*. We describe the two types of projections restricted to Θ^v , as well as the target spaces of said projections, in §3.1 and §3.3, where we also explain some of their basic features. Next we define the window map and prove what is needed from it. Finally, we define Π^v and prove the key properties of the associated projections.

3.1. Quasimorphism distances. For each $v \in \mathcal{V}$, we will use ideas from work-in-progress of the fourth author with Hagen, Russell, and Spriano [HRSS21] to define a map

$$\lambda^v: \Theta^v \rightarrow \mathcal{K}^v,$$

where \mathcal{K}^v is a discrete set quasi-isometric to \mathbb{R} . The key properties of this map are given by the next proposition. We note that the proposition and Lemma 3.6 can be used as black-boxes (in particular, the definitions of λ^v and \mathcal{K}^v are never used after we prove those results).

Proposition 3.1. *There exists $K_1 > 0$ such that, for each $v \in T_\alpha^{(0)} \subset \mathcal{V}$, there exist a space \mathcal{K}^v that is (K_1, K_1) -quasi-isometric to \mathbb{R} and a map $\lambda^v: \Theta^v \rightarrow \mathcal{K}^v$ satisfying the following properties:*

- (1) λ^v is K_1 -coarsely Lipschitz with respect to the path metric on Θ^v .
- (2) For any $x \in \partial\Theta^v$, if $\ell_{x,\alpha} = D_x \cap \partial\mathcal{B}_\alpha$ then $\lambda^v(\ell_{x,\alpha})$ is a set of diameter bounded by K_1 .
- (3) For any $v, w \in \mathcal{V}$ with $d_{\text{tree}}(v, w) = 1$, $\lambda^v \times \lambda^w: \Theta^v \cap \Theta^w \rightarrow \mathcal{K}^v \times \mathcal{K}^w$ is a K_1 -coarsely surjective (K_1, K_1) -quasi-isometry with respect to the induced path metric on the domain.
- (4) (Equivariance) For any $g \in \Gamma$ and $v \in T_\alpha^{(0)}$ there is an isometry $g: \mathcal{K}^v \rightarrow \mathcal{K}^{g^v}$ and for all $x \in \Theta^v$ we have $\lambda^{g^v}(gx) = g\lambda^v(x)$.

The sets $\ell_{x,\alpha}$ in item (2) are certain lines whose significance is explained below.

Remark 3.2. An earlier version of this paper used work of Kapovich and Leeb to construct the spaces \mathcal{K}^v and maps λ^v , resulting in a weaker version of this proposition which did not include the last, equivariance condition. Consequently Γ could only be shown to be an HHS, rather than an HHG. The ideas from [HRSS21] were crucial in this extension.

To explain the proof of the proposition, it is useful to review some background on graph manifolds, which we do now.

Graph manifolds and trees. Recall that a graph manifold is a 3–manifold that contains a canonical finite union of tori (up to isotopy), so that cutting along the tori produces a disjoint union of Seifert fibered 3–manifolds, called the *Seifert pieces*. Seifert fibered 3–manifolds are compact 3–manifolds foliated by *circle leaves*; see [JS79].

The universal cover of a graph manifold decomposes into a union of universal covers of the Seifert pieces glued together along 2–planes (covering the tori). The decomposition is dual to a tree, and the universal covers of the Seifert pieces are *the vertex spaces*. For any Seifert fibered space, its universal cover is foliated by lines, the lift of the foliation by circles, and we refer to the leaves simply as *lines* in the universal cover.

Horocycles and bundles. Next we describe the specific graph manifolds that are relevant for our purposes.

Let $G_\alpha < G$ denote the stabilizer of B_α , for each $\alpha \in \mathcal{P}$. This has a finite index cyclic subgroup G_α^0 generated by a multitwist, $\langle \tau_\alpha \rangle = G_\alpha^0 < G_\alpha$; see e.g. [DDLS21, §2.9]. The preimage of G_α in Γ is the $\pi_1 S$ –extension group Γ_α of G_α , and we likewise denote by $\Gamma_\alpha^0 < \Gamma_\alpha$ the extension group of G_α^0 . The action of Γ_α on $\partial\mathcal{B}_\alpha$ is cocompact, and $\partial\mathcal{B}_\alpha/\Gamma_\alpha$ has a finite sheeted (orbifold) covering by $\partial\mathcal{B}_\alpha/\Gamma_\alpha^0$, which is the graph manifold mentioned in the introduction.

Consider the surface S with the flat metric q_{X_α} , so that $(S, X_\alpha, q_{X_\alpha}) = E_{X_\alpha}/\pi_1 S$. The multitwist τ_α is an affine map that preserves the cylinders in direction α , acting as a power of a Dehn twist in each cylinder and as the identity on their boundaries. The union of the boundaries of the cylinders are spines (deformation retracts) for the subsurfaces that are the complements of the twisting curves (core curves of the cylinders). Consequently, τ_α is the identity on these spines. The homeomorphism τ_α induces a homeomorphism on the subsurface obtained by cutting open S along a core curve of each cylinder. Each such induced homeomorphism is the identity on the corresponding spine, and is thus isotopic to the identity relative to the spine. The mapping torus of each subsurface is a product of the subsurface times a circle, and embeds in the mapping torus $\partial\mathcal{B}_\alpha/\Gamma_\alpha^0$ of τ_α . These sub-mapping tori are the Seifert pieces for the graph manifold structure on $\partial\mathcal{B}_\alpha/\Gamma_\alpha^0$.

The lifted graph manifold decomposition of $\partial\mathcal{B}_\alpha$ corresponds to T_α . That is, for each $v \in T_\alpha^{(0)}$, there is a vertex space contained in Θ^v and containing θ^v . In fact, with respect to the covering group, Θ^v is an invariant, bounded neighborhood of the vertex space and θ^v is an equivariant deformation retraction of that space. We let $\Gamma^v < \Gamma_\alpha$ denote the stabilizer of Θ^v in Γ_α and $\Gamma^{v0} < \Gamma_\alpha^0$ the stabilizer in Γ_α^0 . The suspension flow on the mapping torus $\partial\mathcal{B}_\alpha/\Gamma_\alpha^0$ restricted to each spine defines circle leaves of the corresponding Seifert piece; that is, flow lines through any point on the spine are precisely the circle leaves. In the universal covering $\partial\mathcal{B}_\alpha$, the lifted flowline through a point $x \in \partial\mathcal{B}_\alpha$ is a lifted horocycle, $\ell_{x,\alpha} = D_x \cap \partial\mathcal{B}_\alpha$. Thus, for any vertex v and any $x \in \theta^v$, $\ell_{x,\alpha}$ is a line for the vertex space corresponding to v . We note that not only does Γ^{v0} preserve this set of lines, but so does Γ^v .

For any $x \in \theta^v$, the stabilizer in Γ^{v0} of $\ell_{x,\alpha}$ is generated by a lift g_v of τ_α . Therefore, the quotient Θ^v/Γ^{v0} is homeomorphic to a product, $\Theta_X^v/\pi_1 S^v \times S^1$, where $\pi_1 S^v$ is the stabilizer of v in $\pi_1 S < \Gamma$ and $X \subset \partial\mathcal{B}_\alpha$ is any point. Indeed, there is a deformation retraction to $\theta^v/\Gamma^{v0} = \theta_X^v/\pi_1 S^v \times S^1$. If we do not care about the particular point X over which we take the fiber, we simply write S^v for the surface $\Theta_X^v/\pi_1 S^v$, so that $\Theta^v/\Gamma^{v0} \cong S^v \times S^1$. Since E_X is a copy of the universal

cover of S , we can consider S^v as a subsurface of S (embedded on the interior) and $\pi_1 S^v$ is its fundamental group inside $\pi_1 S$ (up to conjugacy).

The product structure $S^v \times S^1 \cong \Theta^v/\Gamma^{v0}$ can be chosen so that $\Theta^v/\Gamma^{v0} \rightarrow \Theta^v/\Gamma$ is an orbifold cover sending circles to circles making Θ^v/Γ^v into a Seifert fibered *orbifold* (some of the Seifert fibers may be part of the orbifold locus) that also (orbifold)-fibers over the circle (with finite order monodromy). Write $\Theta^v/\Gamma^v \rightarrow \mathcal{O}^v$ for the Seifert fibration to the quotient 2-orbifold. Further write

$$\nu^v: \Gamma^v \rightarrow \pi_1^{orb}(\mathcal{O}^v)$$

for the induced homomorphism of the Seifert fibration and

$$\phi^v: \Gamma^v \rightarrow \mathbb{Z}$$

for the induced homomorphism from the fibration over the circle. Because g_v acts as translation on the line $\ell_{x,\alpha}$ for $x \in \theta^v$, it represents a loop that traverses a circle in the Seifert fibration, which is thus also a suspension flowline for the fibration over the circle. Thus we have $\nu^v(g_v) = 0$ and $\phi^v(g_v) \neq 0$. To complete the picture, we note that restricting $S^v \times S^1 \cong \Theta^v/\Gamma^{v0} \rightarrow \Theta^v/\Gamma^v \rightarrow \mathcal{O}^v$ to S^v defines an orbifold covering $S^v \rightarrow \mathcal{O}^v$.

Finally, note that for any w adjacent to v in T_α , $\langle g_v \rangle \times \langle g_w \rangle \cong \mathbb{Z}^2$ has finite index in $\Gamma^v \cap \Gamma^w$. Viewing $\langle g_w \rangle < \Gamma^v$, we note that $\nu^v|_{\langle g_w \rangle}$ is an isomorphism onto an infinite cyclic subgroup of $\pi_1^{orb} \mathcal{O}^v$. In fact, the image $\nu^v(\langle g_w \rangle)$ is (a conjugate of a power of) the fundamental group of a boundary component of \mathcal{O}^v .

Remark 3.3. One caveat about the lines for the vertex spaces: flowlines through points *not* on a spine are *not* lines of any vertex space. In fact, they are not even uniformly close to lines for any vertex space.

Constructing the map. Here we define \mathcal{K}^v and λ^v and prove the main properties we will need about them. We require a little more set up first. We choose representatives of the Γ -orbits of vertices, $\mathcal{V}_0 = \{v_1, \dots, v_k\} \subset \mathcal{V}$. For each $v \in \mathcal{V}_0$, choose a fundamental domain Δ^v for the action of Γ^v on Θ^v . We assume that Δ^v has compact, connected closure, that $g\Delta^v \cap \Delta^v = \emptyset$ for all $g \in \Gamma^v \setminus \{1\}$, and that $\bigcup_{g \in \Gamma^v} g \cdot \Delta^v = \Theta^v$. The set

$$(1) \quad \{g \in \Gamma^v \mid g\bar{\Delta}^v \cap \bar{\Delta}^v \neq \emptyset\}$$

is a finite generating set for Γ^v . The Γ^v -translates of Δ^v define a tiling of Θ^v , and the map sending every point of $g\Delta^v$ to $g \in \Gamma^v$ is a quasi-isometry by the Milnor-Schwarz Lemma. We denote this map as $\tilde{\lambda}^v: \Theta^v \rightarrow \Gamma^v$.

We note that any word metric on Γ^v defines a “word metric” on each coset $g\Gamma^v$, for $g \in \Gamma$ (elements are distance 1 if they differ by right multiplication by an element of the generating set). We can push the tiling forward by g to a $\Gamma^{g^v} = g\Gamma^v g^{-1}$ -invariant tiling of Θ^{g^v} (if $g \in \Gamma^v$, this is precisely the given tiling of Θ^v). For any element $g' \in \Gamma^{g^v}$, the map that sends every point in $g'\Delta^v$ to g' defines a quasi-isometry $\tilde{\lambda}^{g^v}: \Theta^{g^v} \rightarrow g'\Gamma^{g^v}$ which is Γ^{g^v} -equivariant, with the same quasi-isometry constants. If $g' \in \Gamma$ and $x' \in \Delta^v$, then for all $g \in \Gamma$

$$\tilde{\lambda}^{g^v}(gg'x') = gg' = g\tilde{\lambda}^{g^v}(g'x').$$

On the other hand, any $w \in \Gamma \cdot v$ and $x \in \Theta^w$ have the form $w = g'v$ and $x = g'x'$ for some $g' \in \Gamma$ and $x' \in \Delta^v$. Thus, for any $g \in \Gamma$, the equation above becomes

$$(2) \quad \tilde{\lambda}^{g^w}(gx) = \tilde{g}\lambda^w(x)$$

Having carried out the construction above for each $v \in \mathcal{V}_0$ and each vertex in its orbit, we have maps $\tilde{\lambda}^w$ from Θ^w to a coset of a vertex stabilizer from \mathcal{V}_0 for every $w \in \mathcal{V}$, so that equation (2) holds for every $x \in \Theta^w$, and $g \in \Gamma$.

Next, recall that a *homogeneous quasimorphism (with deficiency D)* from a group H to \mathbb{R} is a map

$$\psi: H \rightarrow \mathbb{R}$$

such that for all $h, h_1, h_2 \in H$ and $n \in \mathbb{Z}$ we have $\psi(h^n) = n\psi(h)$ and

$$|\psi(h_1 h_2) - \psi(h_1) - \psi(h_2)| \leq D.$$

Lemma 3.4. *For any $v \in \mathcal{V}$, there is a homogeneous quasimorphism $\psi^v: \Gamma^v \rightarrow \mathbb{R}$ such that $\psi^v(\langle g_v \rangle)$ is unbounded, and $\psi^v(g_w) = 0$ for any adjacent vertex $w \in \mathcal{V}$.*

Proof. Let w_1, \dots, w_r be Γ^v -orbit representatives of the vertices adjacent to v . Here r is the number of boundary components of \mathcal{O}^v , so that $\nu^v(g_{w_1}), \dots, \nu^v(g_{w_r})$ are peripheral loops around the r distinct boundary components of \mathcal{O}^v . Since $\pi_1^{orb} \mathcal{O}^v$ is the fundamental group of a hyperbolic 2-orbifold with non-empty boundary, appealing to [HO13, Theorem 4.2], which applies to $\pi_1^{orb} \mathcal{O}^v$ and its subgroups $\langle \nu^v(g_{w_i}) \rangle$ in view of [DGO17, Corollary 6.6, Theorem 6.8], one can find a homogeneous quasimorphism $\eta_i: \pi_1^{orb} \mathcal{O}^v \rightarrow \mathbb{R}$, for $i = 1, \dots, r$, such that $\eta_i(\nu^v(g_{w_i})) = 1$ and $\eta_i(\nu^v(g_{w_j})) = 0$ for $j \neq i$. (The construction of Epstein–Fujiwara [EF97] should also be applicable to construct such quasimorphisms). Set $s_0 = 1/\phi^v(g_v)$, and for each $i = 1, \dots, r$, set $s_i = s_0 \phi^v(g_{w_i})$, and then define

$$\psi^v = s_0 \phi^v - \sum_{i=1}^r s_i \eta_i \circ \nu^v.$$

As a linear combination of homogeneous quasimorphisms, ψ^v is a homogeneous quasimorphism. Since $g_v \in \ker(\nu^v)$, it follows that $\eta_i \circ \nu^v(g_v) = 0$ for all i , hence $\psi^v(g_v) = s_0 \phi^v(g_v) = \phi^v(g_v)/\phi^v(g_v) = 1$. On the other hand, for any $j = 1, \dots, r$ we have

$$\psi^v(g_{w_j}) = s_0 \phi^v(g_{w_j}) - \sum_{i=1}^r s_i \eta_i(\nu^v(g_j)) = s_0 \phi^v(g_{w_j}) - \sum_{i=1}^r s_0 \phi^v(g_{w_i}) \delta_{ij} = 0,$$

proving the lemma. \square

According to [ABO19, Lemma 4.15], there is an (infinite) generating set for Γ^v so that with respect to the resulting word metric, the quasimorphism $\psi^v: \Gamma^v \rightarrow \mathbb{R}$ from Lemma 3.4 is a quasi-isometry. For $v \in \mathcal{V}_0$, define $\mathcal{K}^v = \Gamma^v$ with this choice of word metric and let

$$\lambda^v: \Theta^v \rightarrow \mathcal{K}^v$$

simply be the map $\tilde{\lambda}^v$ (followed by the identification of Γ^v with \mathcal{K}^v). For any $g \in \Gamma$, define \mathcal{K}^{gv} to be the coset $g\Gamma^v$ with this generating set so that $\tilde{\lambda}^{gv}$ defines a map

$$\lambda^{gv}: \Theta^{gv} \rightarrow \mathcal{K}^{gv}.$$

Carrying this out for every $v \in \mathcal{V}_0$, (2) implies

$$(3) \quad \lambda^{gw}(gx) = g\lambda^w(x)$$

for all $w \in \mathcal{V}$ and $x \in \Theta^w$, and $g \in \Gamma$.

Before we proceed to the proof of Proposition 3.1, observe that $\Gamma^{gv} = g\Gamma^v g^{-1}$ acts isometrically on $g\Gamma^v$ with respect to any generating set, and thus we can use this to define a generating set for the conjugate so that (any) orbit map is an isometry; in

fact, this will just be a conjugate of the generating set for Γ^v . In particular, when convenient we will identify \mathcal{K}^{g^v} isometrically with the conjugate $g\Gamma^v g^{-1}$ via such an orbit map. Conjugating the quasimorphisms ψ^v from the lemma, for $v \in \mathcal{V}_0$, we obtain uniform quasi-isometries

$$(4) \quad \psi^w: \mathcal{K}^w \rightarrow \mathbb{R}$$

for all $w \in \mathcal{V}$, which for an appropriate choice of identification of \mathcal{K}^w with a conjugate of some Γ^v , $v \in \mathcal{V}_0$, is a quasimorphism (with uniformly bounded deficiency).

Proof of Proposition 3.1. From the discussion above and Equation (3), we immediately see that item (4) of the proposition holds.

Next, observe that by adding finitely many generators to the infinite generating set of Γ^{v_0} for any $v_0 \in \mathcal{V}_0$, changes \mathcal{K}^{v_0} by quasi-isometry. On the other hand, the finite generating set described in Equation (1) for $v_0 \in \mathcal{V}_0$ makes $\tilde{\lambda}^{v_0}$ a quasi-isometry. Thus, adding these generators to the infinite generating set does not change the quasi-isometry type of \mathcal{K}^{v_0} , but clearly makes λ^{v_0} coarsely Lipschitz. Therefore, λ^v is uniformly coarsely Lipschitz for all $v \in \mathcal{V}$, and hence item (1) holds for all $v \in \mathcal{V}$.

To prove item (2), let $v \in \mathcal{V}$ and $x \in \partial\Theta^v$. Then $x \in \theta^w$, for some $w \in \mathcal{V}$ adjacent to v . As discussed above, we view \mathcal{K}^v and \mathcal{K}^w as conjugates Γ^v and Γ^w of groups Γ^{v_0} and Γ^{w_0} , respectively, for $v_0, w_0 \in \mathcal{V}_0$, equipped with their conjugated infinite generating sets. Let $\psi^v: \mathcal{K}^v \rightarrow \mathbb{R}$ and $\psi^w: \mathcal{K}^w \rightarrow \mathbb{R}$ be the associated uniform quasi-isometric homogeneous quasimorphisms. The element $g_w \in \Gamma^v$ stabilizes $\ell_{x,\alpha}$ acting by translation on it, and by construction, $\psi^v(g_w) = \psi^v(g_w^n) = 0$ for all $n \in \mathbb{Z}$. It follows that every orbit of $\langle g_w \rangle$ acting on \mathcal{K}^v is uniformly bounded. Indeed, if D is the deficiency of ψ^v , then for any $g \in \mathcal{K}^v$, we have

$$|\psi^v(g_w^n g) - \psi^v(g)| = |\psi^v(g_w^n g) - \psi^v(g) - \psi^v(g_w^n)| \leq D$$

and therefore $g_w^n g$ and g are uniformly bounded distance apart in \mathcal{K}^v (since ψ^v is a uniform quasiisometry).

Now, since $g_w^n v = v$, by item (4) of the proposition we have

$$\lambda^v(g_w^n x) = \lambda^{g_w^n v}(g_w^n x) = g_w^n \lambda^v(x),$$

and since $g_w^n \lambda^v(x)$ is uniformly close to $\lambda^v(x)$, it follows that λ^v sends the $\langle g_w \rangle$ -orbit of x to a uniformly bounded set. Since this orbit is R -dense in $\ell_{x,\alpha}$ for some uniform $R > 0$, and since λ^v is uniformly coarsely Lipschitz (by item (1)) we see that $\lambda^v(\ell_{x,\alpha})$ has uniformly bounded diameter. This proves item (2).

For item (3), we continue with the assumptions on v, w as above. Note that since $\psi^v(g_v^n) = n$, using again the fact that ψ^v is a uniform quasi-isometric homogeneous quasimorphism to \mathbb{R} , it follows that for any $x \in \Theta^v$, the map $n \mapsto \lambda^v(g_v^n x)$ is a uniformly coarsely surjective, uniform quasiisometry $\mathbb{Z} \rightarrow \mathcal{K}^v$. Since every orbit of $\langle g_w \rangle$ on \mathcal{K}^v is uniformly bounded, it follows that for all $n, m \in \mathbb{Z}$, the two points $\lambda^v(g_v^n g_w^m x) = g_w^m \lambda^v(g_v^n x)$ and $\lambda^v(g_v^n x)$ are uniformly close to each other. Likewise, $\lambda^w(g_v^n g_w^m x)$ and $\lambda^w(g_w^m x)$ are also uniformly close to each other. But this means that

$$\lambda^v \times \lambda^w(g_v^n g_w^m x) \text{ and } (\lambda^v(g_v^n x), \lambda^w(g_w^m x))$$

are uniformly close, and thus

$$(n, m) \mapsto \lambda^v \times \lambda^w(g_v^n g_w^m x)$$

is a uniformly coarsely surjective, uniform quasiisometry $\mathbb{Z}^2 \rightarrow \mathcal{K}^v \times \mathcal{K}^w$.

On the other hand, the assignment $(n, m) \mapsto g_v^n g_w^m x$ defines a uniform quasi-isometry $\mathbb{Z}^2 \rightarrow \Theta^v \cap \Theta^w$ since $\langle g_w \rangle \times \langle g_v \rangle \cong \mathbb{Z}^2$ acts cocompactly on $\Theta^v \cap \Theta^w$ (with uniformity coming from the fact that there are only finitely many Γ -orbits of pairs (v, w) of adjacent vertices). Combining these two facts, together with the fact that λ^v and λ^w are uniformly coarsely Lipschitz, it follows that

$$\lambda^v \times \lambda^w: \Theta^v \cap \Theta^w \rightarrow \mathcal{K}^v \times \mathcal{K}^w$$

is a uniformly coarsely surjective, uniform quasiisometry. This proves item (3), and completes the proof of the proposition. \square

3.2. A technical lemma. The goal of this subsection is to prove Lemma 3.6, whose relevance will only be clear in §4. We prove it here since we have now established the setup for its proof.

We recall that for each v , since Θ^v/Γ^v is a Seifert fibered orbifold, we have have a Γ^v -equivariant, uniformly biLipschitz map $\mu^v \times \rho^v: \Theta^v \rightarrow \tilde{S}^v \times \mathbb{R}$, where \tilde{S}^v is the surface-with-boundary Θ_X^v for some $X \in \partial B_\alpha$ (and $\alpha = \alpha(v)$) and the slices $\{x\} \times \mathbb{R}$ (more precisely, the level sets of $(\mu^v)^{-1}(x) \subset \Theta^v$) are lines for Θ^v . These lines project to circle fibers in Θ^v/Γ^v and we may assume they contain all the lines $\ell_{x,v}$ for all $x \in \theta^v$.

Lemma 3.5. *The map $\mu^v \times \lambda^v: \Theta^v \rightarrow \tilde{S}^v \times \mathcal{K}^v$ is a uniform, Γ^v -equivariant quasi-isometry with uniformly dense image. Moreover, the constant K_1 from Proposition 3.1 can be chosen so that for any $v \in \mathcal{V}$ and $s \in \mathcal{K}^v$, the subspace*

$$M(s) = (\lambda^v)^{-1}(N_{K_1}(s)) \subset \Theta^v,$$

has the property that $\mu^v \times \lambda^v(M(s))$ has uniformly bounded Hausdorff distance to the slice $\tilde{S}^v \times \{s\}$, and furthermore $M(s)$ nontrivially intersects every line of Θ^v .

We note that the intersection of $M(s)$ with each line of Θ^v is necessarily a uniformly bounded diameter set by the uniform bounded Hausdorff distance condition.

Proof. All constants will be independent of the specific vertex v , so we drop it from the notation. We write d for all path-metric distances in what follows (the location of points will determine which metric is being used). Products are given the L^1 metric for convenience. We further let K be the maximum of the coarse Lipschitz constants of μ, ρ, λ and the biLipschitz constant of $\mu \times \rho$, and assume, as we may, that $K \geq 2$. From the proof of Proposition 3.1(3), if x is any point of a line of Θ , then $n \mapsto \lambda^v(g_v^n x)$ is a uniformly coarsely surjective, uniform quasi-isometry from \mathbb{Z} to \mathcal{K} . Therefore, $\lambda = \lambda^v$ is a uniformly coarsely surjective, uniform quasi-isometry from any line of Θ to \mathcal{K} . We further assume that the coarse surjectivity constants and quasi-isometry constants are also all taken to be K .

Let $x, y \in \Theta$ be any two points. Since μ and λ are K -coarsely Lipschitz, $\mu \times \lambda$ is $(2K, 2K)$ -coarsely Lipschitz. To prove the required uniform lower bound on $\mu \times \lambda$ -distances, we note that since $\mu \times \rho$ is K -biLipschitz, it suffices to uniformly coarsely bound $d(\mu \times \rho(x), \mu \times \rho(y))$ from above by $d(\mu \times \lambda(x), \mu \times \lambda(y))$. For reasons that will become clear shortly, we observe that

$$\begin{aligned} d(\mu \times \rho(x), \mu \times \rho(y)) &= d(\mu(x), \mu(y)) + d(\rho(x), \rho(y)) \\ (5) \quad &\leq 2 \max \{ K^4 d(\mu(x), \mu(y)), d(\rho(x), \rho(y)) \}. \end{aligned}$$

If the maximum is realized by $K^4 d(\mu(x), \mu(y))$, then note that

$$\begin{aligned} d(\mu \times \rho(x), \mu \times \rho(y)) &\leq 2K^4 d(\mu(x), \mu(y)) \\ &\leq 2K^4 d(\mu(x), \mu(y)) + 2K^4 d(\lambda(x), \lambda(y)) \\ &= 2K^4 d(\mu \times \lambda(x), \mu \times \lambda(y)), \end{aligned}$$

as required.

We are left to consider the case that the maximum in (5) is realized by $d(\rho(x), \rho(y))$, which thus satisfies

$$d(\rho(x), \rho(y)) \geq K^4 d(\mu(x), \mu(y)).$$

Let $z \in \Theta$ be such that $\rho(z) = \rho(x)$ and $\mu(z) = \mu(y)$. Since $\mu(z) = \mu(y)$, z and y lie on a line, and since the restriction of λ to this line is a (K, K) -quasi-isometry, we have

$$d(\lambda(z), \lambda(y)) \geq \frac{1}{K} d(\rho(z), \rho(y)) - K = \frac{1}{K} d(\rho(x), \rho(y)) - K.$$

Since λ is K -coarsely Lipschitz and $\mu \times \rho$ is K -biLipschitz, we have

$$\begin{aligned} d(\lambda(x), \lambda(z)) &\leq K d(x, z) + K \\ &\leq K^2 (d(\mu(x), \mu(z)) + d(\rho(x), \rho(z))) + K \\ &= K^2 d(\mu(x), \mu(y)) + K \\ &\leq K^2 \left(\frac{1}{K^4} d(\rho(x), \rho(y)) \right) + K \\ &= \frac{1}{K^2} d(\rho(x), \rho(y)) + K \end{aligned}$$

Combining the previous two sets of inequalities and the triangle inequality, we have

$$\begin{aligned} d(\lambda(x), \lambda(y)) &\geq d(\lambda(z), \lambda(y)) - d(\lambda(z), \lambda(x)) \\ &\geq \frac{1}{K} d(\rho(x), \rho(y)) - K - \left(\frac{1}{K^2} d(\rho(x), \rho(y)) + K \right) \\ &\geq \frac{K-1}{K^2} d(\rho(x), \rho(y)) - 2K. \end{aligned}$$

Combining this inequality with (5) where we have assumed the maximum is realized by $d(\rho(x), \rho(y))$, we obtain

$$\begin{aligned} d(\mu \times \rho(x), \mu \times \rho(y)) &\leq 2d(\rho(x), \rho(y)) \\ &\leq \frac{2K^2}{K-1} d(\lambda(x), \lambda(y)) + \frac{4K^3}{K-1} \\ &\leq \frac{2K^2}{K-1} d(\mu \times \lambda(x), \mu \times \lambda(y)) + \frac{4K^3}{K-1}. \end{aligned}$$

which provides the required upper bound. This completes the proof of the first claim of the lemma.

For the second claim of the lemma, we now increase K_1 from Proposition 3.1 if necessary, so that $K_1 \geq K$. Observe that

$$(6) \quad (\mu \times \lambda) \circ (\mu \times \rho)^{-1}(x, t) = (x, \lambda((\mu \times \rho)^{-1}(x, t))).$$

That is, $(\mu \times \lambda) \circ (\mu \times \rho)^{-1}$ sends the line $\{x\} \times \mathbb{R}$ to $\{x\} \times \mathcal{K}$, for any $x \in \tilde{S}$. As already noted at the start of the proof, restricting to this line, λ is K -coarsely Lipschitz and K -coarsely onto. Therefore, for any $s \in \mathcal{K}$ and $x \in \tilde{S}$, there exists t so that $\lambda((\mu \times \rho)^{-1}(x, t))$ is within $K_1 \geq K$ of s . Thus, for any line of Θ , the λ -image nontrivially intersects $N_{K_1}(s)$, and hence this line nontrivially intersects $M(s)$. By definition, $\mu \times \lambda$ maps $M(s)$ into $\tilde{S}^v \times N_{K_1}(s)$, and by the previous sentence, every point of $\tilde{S} \times \{s\}$ is within K_1 of some point of $\mu \times \lambda(M(s))$. Thus, $\lambda \times \mu(M(s))$ has Hausdorff distance at most K_1 from $\tilde{S} \times \{s\}$, as required. \square

As mentioned above, the following technical lemma will be needed in §4. In the statement $M(s)$ is as defined by Lemma 3.5.

Lemma 3.6. *There is a function $N = N_{\mathcal{K}}: [0, \infty) \rightarrow [0, \infty)$ with the following property. Suppose that $v, v_1, v_2 \in \mathcal{V}$ are so that $d_{\text{tree}}(v, v_i) = 1$. Then for each $s \in \mathcal{K}^v$ and $t_i \in \mathcal{K}^{v_i}$ we have*

$$\bar{d}(M(s) \cap M(t_1), M(s) \cap M(t_2)) \leq N(\bar{d}(M(t_1), M(t_2))).$$

Proof. It suffices to prove the lemma with \bar{d} replaced by the path metric $d_{\partial\mathcal{B}_\alpha}$ on $\partial\mathcal{B}_\alpha$, since they are uniformly coarsely equivalent. In fact, it will be convenient to consider the path metric d_0 on the union of the three vertex subspaces

$$\Omega = \Theta_v \cup \Theta_{v_1} \cup \Theta_{v_2},$$

which is also uniformly coarsely equivalent since each vertex space uniformly quasi-isometrically embeds in $\partial\mathcal{B}_\alpha$. In this subspace, we will actually prove that the two distances are uniformly comparable.

Now, for each $i = 1, 2$ the uniform quasi-isometry $\mu^{v_i} \times \lambda^{v_i}: \Theta^{v_i} \rightarrow \tilde{S}^{v_i} \times \mathcal{K}^{v_i}$ from Lemma 3.5 maps the space $\Theta^v \cap \Theta^{v_i}$ within bounded Hausdorff distance of a subspace $\partial^v \tilde{S}^{v_i} \times \mathcal{K}^{v_i}$, for a boundary component $\partial^v \tilde{S}^{v_i}$ of \tilde{S}^{v_i} . Let $\eta: \tilde{S}^{v_i} \rightarrow \partial^v \tilde{S}^{v_i}$ be the closest point projection, and then set

$$(\mu^{v_i} \times \lambda^{v_i})^{-1} \circ (\eta \times id) \circ (\mu^{v_i} \times \lambda^{v_i}): \Theta^{v_i} \rightarrow \Theta^{v_i}$$

where $(\mu^{v_i} \times \lambda^{v_i})^{-1}$ is a coarse inverse of $\mu^{v_i} \times \lambda^{v_i}$ with $\mu^{v_i} \circ (\mu^{v_i} \times \lambda^{v_i})^{-1}(x, s) = x$ (c.f. Equation (6)). This map is a uniformly coarsely Lipschitz, coarse retraction of Θ^{v_i} onto $\Theta^{v_i} \cap \Theta^v$. Moreover, this sends $M(t_i)$, which is uniformly close to the $(\mu^{v_i} \times \lambda^{v_i})^{-1}$ -image of $\tilde{S}^{v_i} \times \{t_i\}$, to a uniformly bounded neighborhood of $M(t_i) \cap \Theta^v$. Consequently,

$$(7) \quad d_0(M(t_1), M(t_2)) \asymp d_0(M(t_1) \cap \Theta^v, M(t_2) \cap \Theta^v)$$

with uniform constants.

Next, observe that $M(t_i) \cap \Theta^v \subset \Theta^v \cap \Theta^{v_i} \subset \Theta^v$.

Claim 3.7. *The quasi-isometry $\mu^v \times \lambda^v$ maps $M(t_i) \cap \Theta^v$ within a uniformly bounded Hausdorff distance of the slice $\{z_i\} \times \mathcal{K}^v \subset \tilde{S}^v \times \mathcal{K}^v$, for each $i = 1, 2$, where (z_i, t'_i) is a point in the $\mu^v \times \lambda^v$ -image of $M(s) \cap M(t_i)$.*

Assuming the claim, we note then that

$$\begin{aligned} d_0(M(t_1) \cap \Theta^v, M(t_2) \cap \Theta^v) &\asymp d_{\tilde{S}^v \times \mathcal{K}^v}(\{z_1\} \times \mathcal{K}^v, \{z_2\} \times \mathcal{K}^v) \\ &= d_{\tilde{S}^v}(z_1, z_2) \\ &\asymp d_0(M(s) \cap M(t_1), M(s) \cap M(t_2)) \end{aligned}$$

again with uniform constants. Combining this coarse equation with (7) we get the required uniform estimate

$$d_0(M(s) \cap M(t_1), M(s) \cap M(t_2)) \asymp d_0(M(t_1), M(t_2)).$$

Fix $i = 1$ or 2 and we prove the claim. Since $\lambda^v \times \lambda^{v_i}$ is K_1 -coarsely surjective (Proposition 3.1(3)), there exists some point $y_i \in \Theta^v \cap \Theta^{v_i}$ with $(\lambda^v(y_i), \lambda^{v_i}(y_i))$ within distance K_1 of $(s, t_i) \in \mathcal{K}^v \times \mathcal{K}^{v_i}$. Therefore, $y_i \in M(s) \cap M(t_i)$ and we set $\mu^v \times \lambda^v(y_i) = (z_i, t'_i)$.

Next, we observe that λ^{v_i} is uniformly coarsely constant on any line of Θ^v contained in $\Theta^v \cap \Theta^{v_i}$ by Proposition 3.1(2) and uniformly coarsely Lipschitz by Proposition 3.1(1). Hence, the line

$$(\mu^v)^{-1}(z_i) = (\mu^v \times \rho^v)^{-1}(\{z_i\} \times \mathbb{R})$$

of Θ^v maps under $\lambda^v \times \lambda^{v_i} : \Theta^v \cap \Theta^{v_i} \rightarrow \mathcal{K}^v \times \mathcal{K}^{v_i}$ into a neighborhood of uniformly bounded radius of $\mathcal{K}^v \times \{t_i\}$. Therefore, any point in the image of the line in $\mathcal{K}^v \times \mathcal{K}^{v_i}$ lies uniformly close to a point in $\lambda^v \times \lambda^{v_i}(M(t_i) \cap \Theta^v)$ by Proposition 3.1(3) (which guarantees that any point in $\mathcal{K}^v \times \{t_i\}$ is K_1 -close to a point in the image of the subspace $\Theta^v \cap \Theta^{v_i}$). Therefore any point in the line lies uniformly close to some point in $M(t_i) \cap \Theta^v$ since $\lambda^v \times \lambda^{v_i}$ is a uniform quasi-isometry again by Proposition 3.1(3).

On the other hand, Lemma 3.5 implies $\mu^{v_i} \times \lambda^{v_i}(M(t_i))$ is uniformly bounded Hausdorff distance from the slice $\tilde{S}^{v_i} \times \{t_i\} \subset \tilde{S}^{v_i} \times \mathcal{K}^{v_i}$. Moreover, since $M(t_i)$ meets every line of Θ^{v_i} (Lemma 3.5 again), it follows that

$$\mu^{v_i} \times \lambda^{v_i}(M(t_i) \cap \Theta^v)$$

is uniformly bounded Hausdorff distance to the quasi-line $\partial^v \tilde{S}^{v_i} \times \{t_i\}$ (see the proof of Lemma 3.5). In particular, $M(t_i) \cap \Theta^v$ is itself a uniform quasi-line and consequently lies within a uniformly bounded neighborhood of the line $(\mu^v)^{-1}(z_i)$. Since this line maps within within a uniformly bounded Hausdorff distance of the slice $\{z_i\} \times \mathcal{K}^v$ in $\tilde{S}^v \times \mathcal{K}^v$ by $\mu^v \times \lambda^v$, we see that $M(t_i) \cap \Theta^v$ does as well \square

3.3. Coned-off surfaces. For $v \in \mathcal{V}$, we define Ξ^v to be the graph whose vertices are all $w \in \mathcal{V}$ so that $d_{\text{tree}}(v, w) = 1$, and with edges connecting the pairs w, w' whenever $\Theta^w \cap \Theta^{w'} \neq \emptyset$. As such, vertices $w \in \mathcal{V}$ are in bijective correspondence with the boundary components of Θ^v and there is an ‘‘inclusion’’ map

$$i^v : \partial\Theta^v \rightarrow \Xi^v$$

that sends any point $x \in \partial\Theta^v$ to the vertex w for which $x \in \Theta^w$. In light of the following lemma, we note that we could alternately define the edges of Ξ^v in terms of subspaces lying within bounded distance of each other, and produce a space quasi-isometric to Ξ^v .

Lemma 3.8. *There is a function $N = N_{\Xi} : [0, \infty) \rightarrow [0, \infty)$ so that whenever $v, w_1, w_2 \in \mathcal{V}$ satisfy $\bar{d}(\Theta^{w_1} \cap \Theta^v, \Theta^{w_2} \cap \Theta^v) \leq r$, the sets $\Theta^{w_1} \cap \Theta^v$ and $\Theta^v \cap \Theta^{w_2}$ may be connected via a concatenation of at most $N(r)$ paths, each of which is contained in a set of the form $\Theta^v \cap \Theta^w$.*

Proof. If $\bar{d}(\Theta^{w_1} \cap \Theta^v, \Theta^{w_2} \cap \Theta^v) \leq r$, then there are cone points $p_i \in \Theta^{w_i} \cap \theta^v$ within distance $r + 3M$, where M is the bound on the width of a strip and length of a saddle connection from Lemma 2.3. Since the path metric d_{θ^v} on θ^v is coarsely equivalent to the subspace metric, $d_{\theta^v}(p_1, p_2)$ is bounded in terms of r . The path metric on θ^v is biLipschitz equivalent to the ℓ^1 -metric on the product $\theta_{X_v}^v \times \mathbb{R}$. Since each edge of $\theta_{X_v}^v$ has definite length, there is a path from p_1 to p_2 in θ^v obtained by concatenating boundedly many (in terms of r) paths α_i with $f_{X_v}(\alpha_i)$ a saddle connection in $\theta_{X_v}^v$. Since each α_i is contained in some Θ^w , we are done. \square

Corollary 3.9. *For any $x, y \in \partial\Theta^v$, we have $d_{\Xi^v}(i^v(x), i^v(y)) \leq N(d_{\Theta^v}(x, y)) + 1$.*

Proof. Let $w = i^v(x)$ and $w' = i^v(y)$. Since d_{Θ^v} and \bar{d} are path metrics, we have $\bar{d}(x, y) \leq d_{\Theta^v}(x, y)$. By Lemma 3.8, x and y may be joined by a concatenation $\alpha_1 \cdots \alpha_k$ of $k \leq N(d_{\Theta^v}(x, y)) + 2$ paths α_j each of which lies in some $\Theta^v \cap \Theta^{w_j}$, and where $w = w_1$ and $w' = w_k$. For successive paths α_j, α_{j+1} , the vertices w_j, w_{j+1} are adjacent in Ξ^v by definition. Therefore $d_{\Xi^v}(w, w') \leq k - 1 \leq N(d_{\Theta^v}(x, y)) + 1$. \square

Lemma 3.10. *Each Ξ^v is uniformly quasi-isometric to a tree. In particular, there exists $\delta > 0$ so that each Ξ^v is δ -hyperbolic. Moreover, Ξ^v has at least two points at infinity.*

Proof. We appeal to Proposition 2.5 and show that for any vertices w, w' of Ξ^v there exists a path $\gamma(w, w')$ so that any path from w to w' passes within distance 3 of each vertex of $\gamma(w, w')$.

First, note that Ξ^v is isomorphic to the intersection graph of the collection of strips in Θ_0^v . For each strip we have a vertex, and for each saddle connection of the spine θ_0^v , there is an edge of Ξ^v that connects the vertices corresponding to the strips that contain the saddle connection. For each cone point in the spine θ_0^v , there is also a complete graph on the vertices corresponding to strips that contain this cone point. This accounts for all edges (because intersections of strips either arise along saddle connections or single cone points), and we note that the closure of each edge of the first type separates Ξ^v into two components.

Suppose $w, w' \in \Xi^v$ are two vertices, and let $x, x' \in \theta_0^v$ be points in the (boundaries of the) strips corresponding to w and w' , respectively, that are closest in Θ_0^v , and consider the geodesic in θ_0^v connecting these points, which is a concatenation of saddle connections $\sigma_1 \sigma_2 \cdots \sigma_n$. For each $1 \leq i \leq n$, let w_i^\pm be the vertices corresponding to the two strips A_i^\pm that intersect in the saddle connection σ_i . We can form an edge path $\gamma(w, w')$ in Ξ^v , containing w, w' , and the w_i^\pm as vertices, since $A_i^+ \cap A_{i+1}^+ \neq \emptyset$. Observe that any path from x to x' must pass through the union $A_i^+ \cup A_i^-$, for each i , since x and x' lie in the closures of distinct components of $\Theta_0^v \setminus (A_i^+ \cup A_i^-)$.

Now let $w = w_0, w_1, \dots, w_k = w'$ be the vertices of an edge path connecting w to w' in Ξ^v . For any points in the strips corresponding to w and w' , respectively, it is easy to construct a path in Θ_0^v between these points that decomposes as a concatenation $\nu_1 \nu_2 \cdots \nu_k$ so that ν_j is contained entirely in the strip corresponding to w_j . From the previous paragraph, this path must pass through $A_i^+ \cup A_i^-$, for each $i = 1, \dots, k$. It follows that for each $1 \leq i \leq k$, the edge path must meet the union of the stars $\text{star}(w_i^+) \cup \text{star}(w_i^-)$. Since these stars intersect, their union has diameter at most 3, and we are done.

We now show that Ξ^v contains a quasi-geodesic line. Consider strips A_i of Θ_0^v , for $i \in \mathbb{Z}$, such that for all i we have

- A_i and A_{i+1} share a saddle connection;
- A_{i-1} and A_{i+2} lie on distinct components of the complement of the interior of $A_i \cup A_{i+1}$ in Θ_0^v .

The A_i give a bi-infinite path in Ξ^v , and we now show that this path is a quasi-geodesic. Fix integers m, n and consider a geodesic γ in Ξ^v from A_m to A_n (where we think of the strips themselves as vertices of Ξ^v for convenience). Then for each $m < k < n - 1$ we have that γ needs to contain a vertex $v(k)$ which, regarded as a strip, intersects A_k or A_{k+1} . Indeed, the interior of $A_k \cup A_{k+1}$ separates A_m from A_n , and the sequence of vertices of γ corresponds to a connected union of

strips containing A_m and A_n . Moreover, there is no strip intersecting both A_k and $A_{k'}$ if $|k - k'| \geq 3$, and in particular we have $v(k) \neq v(k')$ if $|k - k'| \geq 4$. These observations imply that γ contains at least $\lfloor (n - m - 2)/4 \rfloor$ vertices, so that geodesics connecting A_m to A_n have length comparable to $n - m$, and the A_i form a quasi-geodesic line as required. \square

3.4. Windows and bridges. For $v \in \mathcal{V}$, consider the set $\bar{\Sigma}^v$ of points in $\bar{\Sigma}$ that are inside some v -spine, as well as those points $\bar{\Sigma}^{\neq v}$ that are outside every v -spine:

$$\bar{\Sigma}^v = \bigcup_{X \in \bar{D}} (\theta_X^v \cap \Sigma) \quad \text{and} \quad \bar{\Sigma}^{\neq v} = \bar{\Sigma} \setminus \bar{\Sigma}^v.$$

For each $Y \in \bar{D}$ we now define a *window map* $\Pi_Y^v: \bar{\Sigma} \rightarrow \mathbb{P}(\partial\Theta_Y^v)$ from cone points to the set of subsets of the boundary $\partial\Theta_Y^v \subset \Theta_Y^v$. There are two cases. Firstly,

$$\text{for } x \in \bar{\Sigma}^{\neq v}, \quad \Pi_Y^v(x) = \{z \in \partial\Theta_Y^v \mid [f_Y(x), z] \cap \Theta_Y^v = \{z\}\}.$$

In words, $\Pi_Y^v(x)$ is the union of entrance points in Θ_Y^v of any flat geodesic in E_Y from $f_Y(x)$ to Θ_Y^v (basically the closest point projection in E_Y), and we call it the *window* for x in Θ_Y^v . Observe that for any $X, Y \in \bar{D}$ we have $f_{X,Y}(\Pi_Y^v(x)) = \Pi_X^v(x)$. The second case, that of $\bar{\Sigma}^v$, is handled slightly differently:

$$\begin{aligned} \text{for } x \in \bar{\Sigma}^v, \quad \Pi_Y^v(x) &= f_{Y,X}(\Pi_X^v(x)), \quad \text{where } X = c_v(\pi(x)) \in \partial B_v \\ \text{and } \Pi_X^v(x) &= \{z \in \partial\Theta_X^v \mid z \text{ is a closest cone point to } f_X(x) \in \theta_X^v\}. \end{aligned}$$

Thus affine invariance $\Pi_Z^v(x) = f_{Z,Y}(\Pi_Y^v(x))$ is built directly into the definition.

Now for any $Y \in \bar{D}$ and $x, y \in \bar{\Sigma}$, we define

$$d_Y^v(x, y) = \text{diam}(\Pi_Y^v(x) \cup \Pi_Y^v(y)),$$

where the distance is computed in the path metric on E_Y (or equivalently on Θ_Y^v). Finally, we extend window maps to arbitrary subsets by declaring

$$\text{for } U \subset \bar{E}, \quad \Pi_Y^v(U) = \Pi_Y^v(U \cap \bar{\Sigma}) = \bigcup_{x \in \bar{\Sigma} \cap U} \Pi_Y^v(x)$$

which has the same effect as defining $\Pi_Y^v(x) = \emptyset \subset \partial\Theta_Y^v$ for $x \notin \bar{\Sigma}$.

Lemma 3.11. *If $x, y \in \bar{\Sigma}^{\neq v}$ satisfy $f(x) = f(y)$, then $\Pi_Y^v(x) = \Pi_Y^v(y)$ for all $Y \in \bar{D}$.*

Proof. This is immediate since $\Pi_Y^v(x)$ is defined just in terms of $f_Y(x) = f_Y(y)$. \square

The following gives a counterpoint to Lemma 3.11 for points in $\bar{\Sigma}^v$.

Lemma 3.12. *There exists $K_2 > 0$ such that for any $v \in \mathcal{V}$ the following holds: If $x, y \in \bar{\Sigma}^v$ satisfy $f(x) = f(y)$ and either*

- (1) *x and y are connected by a horizontal geodesic of length ≤ 1 , or*
- (2) *x and y are contained in $\partial\mathcal{B}_w$ for some $w \in \mathcal{V}$ with $\alpha(w) \neq \alpha(v)$,*

then $d_X^v(x, y) \leq K_2$, where $X = c_v(\pi(x))$.

Proof. Set $X = c_v(\pi(x))$ and $Y = c_v(\pi(y))$. Since $c_v: D \rightarrow \partial B_v$ is 1-Lipschitz and $\text{diam}(c_v(\partial B_w))$ is uniformly bounded for all such w , either condition (1) or (2) gives a uniform bound $K > 0$ on the distance between X and Y . Hence $f_{X,Y}$ is e^K -biLipschitz. The distance between $f_Y(y) \in \theta_Y^v$ and its closest cone points $\Pi_Y^v(y)$ in $\partial\Theta_Y^v$ is also uniformly bounded by $2M$, by Lemma 2.3. The same holds for the distance between $f_X(x)$ and $\Pi_X^v(x)$. It follows that $\Pi_X^v(y) = f_{X,Y}(\Pi_Y^v(y))$




FIGURE 2. A window

lies within distance $2e^K M$ of $f_X(x) = f_{X,Y}(f_Y(y))$ and hence within distance $2e^K M + 2M$ of $\Pi_X^v(x)$. \square

The next lemma explains that the image of Π_Y^v is not so far from being a point.

Lemma 3.13 (Window lemma). *For any $v \in \mathcal{V}$, $Y \in \bar{D}$, and $x \in \bar{\Sigma}^{\neq v}$, the window $\Pi_Y^v(x) \subset \partial\Theta_Y^v$ is either a cone point or a single saddle connection.*

Proof. If $f_Y(x) \in \partial\Theta_Y^v$ then clearly $\Pi_Y^v(x)$ is the cone point $f_Y(x)$ itself. So suppose $f_Y(x) \notin \Theta_Y^v$ and let ℓ be the component of $\partial\Theta_Y^v$ separating $f_Y(x)$ from θ_Y^v in E_Y , so that $\Pi_Y^v(x) \subset \ell$. Take any flat geodesic $[f_Y(x), z]$ in E_Y from $f_Y(x)$ to a cone point $z \in \ell$. The geodesic $[f_Y(x), z]$ first meets ℓ in some cone point p . Since the total cone angle at p is at least 3π and the angle at p along the side of ℓ containing Θ^v is exactly π , the last saddle connection δ in the geodesic $[f_Y(x), p] \subset [f_Y(x), z]$ must make an angle of at least π with ℓ on one side. Hence any geodesic from $f_Y(x)$ to a cone point on that side of p must pass through p .

If both angles between δ and ℓ at p are at least π , then any geodesic from $f_Y(x)$ to ℓ passes through p . Hence p is the unique point in $\partial\Theta_Y^v$ closest to $f_Y(x)$ and $\Pi_Y^v(x) = \{p\}$ is a cone point as required. Otherwise, consider the flat geodesic from $f_Y(x)$ to the adjacent cone point p' on the other side of p along ℓ . The last saddle connection of this geodesic must also make an angle with ℓ of at least π on one side. This cannot be the side containing p , or else the geodesic from $f_Y(x)$ to p would pass through p' contradicting our choice of p . Hence any geodesic from $f_Y(x)$ to a cone point on the opposite side of p' must pass through p' . Therefore $\Pi_Y^v(x)$ is the saddle connection between p and p' , and we are done. See Figure 2. \square

The following lemma gives us partial control over the window for points in adjacent vertex spaces in the same Bass-Serre tree.

Lemma 3.14 (Bridge lemma). *For any $v, w \in \mathcal{V}$ with $d_{\text{tree}}(v, w) = 1$, any $Y \in \bar{D}$, and any component $U \subset E_Y \setminus \theta_Y^w$ not containing θ_Y^v , there exists a (possibly degenerate) saddle connection $\delta_U \subset \partial\Theta_Y^v$ with the following property: Every $x \in \bar{\Sigma}$ with $f_Y(x) \in \bar{U}$ satisfies $\Pi_Y^v(x) \subset \delta_U$.*

We call δ_U the *bridge* for U in E_Y . It is clear from the construction in the proof below that $f_{Z,Y}(\delta_U)$ is the bridge for $f_{Z,Y}(U)$, for any $Z \in \bar{D}$.

Proof. Let U be as in the statement and W be the component of $E_Y \setminus \theta_Y^w$ containing θ_Y^v . Let $\gamma_U = \bar{U} \cap \theta_Y^w$ and $\gamma_W = \bar{W} \cap \theta_Y^w \subset \partial\Theta_Y^v$, which are both bi-infinite flat geodesics in θ_Y^w .

If $\gamma_U \cap \gamma_W = \emptyset$, then there is a unique geodesic between them in θ_Y^w , and we take δ_U to be the endpoint of this geodesic which lies along γ_W . On the other hand, if $\gamma_U \cap \gamma_W \neq \emptyset$, then their intersection is contained in the boundary of a strip along θ_Y^v and another along θ_Y^w . Two distinct strips in the same direction that intersect do so in either a single point or a single saddle connection, and hence $\gamma_U \cap \gamma_W$ is a point or single saddle connection, and we call δ_U . See Figure 3.




FIGURE 3. Bridges, and the proof of Lemma 3.14.

Now consider any point $x \in \bar{\Sigma}$ with $f_Y(x) \in \bar{U}$. Observe that $x \in \bar{\Sigma}^{\neq v}$ so that $\Pi_Y^v(x)$ falls under the first case of the window definition. Further, any flat geodesic from $f_Y(x)$ to Θ_Y^v must pass through both γ_U and γ_W , and hence must pass through $\delta_U \subset \partial\Theta_Y^v$. It follows that $\Pi_Y^v(x) \subset \delta_U$, as required. \square

The following is an easy consequence of the previous lemma.

Corollary 3.15. *For any $v, w \in \mathcal{V}$ with $2 \leq d_{\text{tree}}(v, w) < \infty$ and $Y \in \bar{D}$, there exists a (possibly degenerate) saddle connection $\delta_Y^v(w) \subset \partial\Theta_Y^v$ so that if $x \in \bar{\Sigma}$ has $f_Y(x) \in \Theta_Y^w$, then $\Pi_Y^v(x) \subset \delta_Y^v(w)$. In particular, $\Pi_Y^v(\Theta^w) \subset \delta_Y^v(w)$.*

Proof. There exists $u \in T_{\alpha(v)}^{(0)}$ between v and w with $d_{\text{tree}}(u, v) = 1$, and a component $U \subset E_Y \setminus \theta_Y^u$ whose closure contains Θ_Y^w . Setting $\delta_Y^v(w) = \delta_U$ and applying Lemma 3.14 completes the proof. \square

The next corollary is similar.

Corollary 3.16. *If $v, w \in \mathcal{V}$ satisfy $d_{\text{tree}}(v, w) = \infty$ and $\theta_0^v \cap \theta_0^w = \emptyset$, then for each $Y \in \bar{D}$ there is a connected union $\delta_Y^v(w) \subset \partial\Theta_Y^v$ of at most two saddle connections such that $\Pi_Y^v(x) \subset \delta_Y^v(w)$ for all $x \in \bar{\Sigma}^w$. In particular $\Pi_Y^v(\theta^w) \subset \delta_Y^v(w)$.*

Proof. From the hypotheses, θ_Y^w is contained in some component $W \subset E_Y \setminus \theta_Y^v$. Let $u \in \mathcal{V}$ be such that $d_{\text{tree}}(v, u) = 1$ and $\theta_Y^u \subset W$. Since $\alpha(u) = \alpha(v) \neq \alpha(w)$, θ_Y^u and θ_Y^w can intersect in at most one point. If $\theta_Y^u \cap \theta_Y^w = \emptyset$, then θ_Y^w is contained in a component U of $E_Y \setminus \theta_Y^u$ disjoint from θ_Y^v ; thus $\Pi_Y^v(\theta_Y^w)$ is contained in the bridge δ_U for U by Lemma 3.14. Otherwise $\theta_Y^u \cap \theta_Y^w \neq \emptyset$, and we claim there is a cone point $p \in \theta_Y^u$ such that $p \in \bar{U}$ for every component U of $E_Y \setminus \theta_Y^u$ that intersects θ_Y^w .

Indeed, if $\theta_Y^v \cap \theta_Y^w$ is a cone point, we take p to be this intersection point, and if not $\theta_Y^v \cap \theta_Y^w$ is an interior point of a saddle connection of θ_Y^v and we may take p to be either of its endpoints. For each component $U \subset E_Y \setminus \theta_Y^v$ that intersects θ_Y^w we then have $\Pi_Y^v(p) \subset \delta_U$, where δ_U is the bridge for U . Since $f_Y(\bar{\Sigma}^w) \subset \theta_Y^w$ is contained in the union of the closures of such U , it follows that $\Pi_Y^v(\bar{\Sigma}^w)$ is contained in a union of saddle connections along $\partial\Theta_Y^v$, all of which contain $\Pi_Y^v(p)$, and hence is a connected union of at most two saddle connections. This completes the proof. \square

The final case to consider is that of spines in different directions that intersect:

Lemma 3.17. *There exists $K_3 > 0$ such that if $v, w \in \mathcal{V}$ with $d_{\text{tree}}(v, w) = \infty$ and $\theta_0^v \cap \theta_0^w \neq \emptyset$, then $\text{diam}(\Pi_Y^v(\theta^w)) < K_3$ for all $Y \in c_v(\partial B_w)$.*




FIGURE 4. The proof of Lemma 3.17

Proof. Let $x_0 = \theta_Y^v \cap \theta_Y^w$ be the unique intersection point of the spines. Let W_0 be the smallest subgraph of θ_Y^w containing $\theta_Y^w \cap \Theta_Y^v$ and let W_1, \dots, W_k be the closures of the components of $\theta_Y^w \setminus W_0$, so that $\theta_Y^w = W_0 \cup \dots \cup W_k$. See Figure 4. For $1 \leq i \leq k$, let $x_i \in W_i$ be the closest (cone) point to x_0 . Then define p_i to be the intersection of the geodesic $[x_0, x_i] \subset \theta_Y^w$ with $\partial\Theta_Y^v$ and let $\delta_i \subset \partial\Theta_Y^v$ be the segment consisting of the (1 or 2) saddle connections along $\partial\Theta_Y^v$ that contain p_i . Define $\delta_Y^v(w) = \delta_1 \cup \dots \cup \delta_k$, and note that $\delta_X^v(w) = f_{X,Y}(\delta_Y^v(w))$ and $\theta_Y^w \cap \partial\Theta_Y^v \subset \delta_Y^v(w)$ by construction.

For any $x \in \bar{\Sigma}$ with $f_Y(x) \in W_i$, where $i = 1, \dots, k$, the flat geodesic from $f_Y(x)$ to x_0 first intersects Θ_Y^v at p_i ; therefore $p_i \in \Pi_Y^v(x)$ by definition of the window. By Lemma 3.13, it follows that $\Pi_Y^v(x) \subset \delta_i$. The union $\cup_{i=1}^k W_i$ contains every cone point of θ_Y^w except possibly x_0 . Thus we have proven $\Pi_Y^v(\theta^w \cap \bar{\Sigma}^{\neq v}) = \Pi_Y^v(\theta_Y^w \cap \bar{\Sigma}^{\neq v}) \subset \delta_Y^v(w)$.

Let $Z \in D$ be the closest point on ∂B_v to ∂B_w , thus $Z \in \partial B_v$ lies on the unique hyperbolic geodesic that intersects ∂B_v and ∂B_w orthogonally. In E_Z the directions $\alpha(v)$ and $\alpha(w)$ are perpendicular. Therefore the cone points of $\partial\Theta_Z^v$ that are closest to $\theta_Z^v \cap \theta_Z^w$ all lie in $\delta_Z^v(w)$, since they must be endpoints of saddle connections along

$\partial\Theta_Z^v$ that intersect θ_Z^w . More generally, for any $Y \in c_v(\partial B_w)$, the directions $\alpha(v)$ and $\alpha(w)$ are nearly perpendicular and thus we have

$$\Pi_Y^v(\theta^w \cap \bar{\Sigma}^{\xi^v}) \subset \delta_Y^v(w) \subset B_K(\theta_Y^v \cap \theta_Y^w).$$

for some uniform constant $K > 0$ that depends only on the length of $c_v(\partial B_w)$ and the maximum over ∂B_v of the length/width of any saddle connection/strip in the $\alpha(v)$ direction (Lemma 2.3). Now, for any $x \in \theta^w \cap \bar{\Sigma}^v$, the point $X = c_v(\pi(x))$ lies in $c_v(\partial B_w)$ and we have that $f_X(x) = \theta_X^v \cap \theta_X^w$. The above equation shows there are cone points of $\partial\Theta_X^v$ within K of $f_X(x)$; hence $\Pi_X^v(x) = \Pi_X^v(f_X(x))$ lies in $B_K(\theta_X^v \cap \theta_X^w)$ by definition. Using the fact that the length of $c_v(\partial B_w)$ is uniformly bounded, we see that the map $f_{Y,X}$ is uniformly biLipschitz and therefore that $\Pi_Y^v(x) = f_{Y,X}(\Pi_X^v(x))$ lies within bounded distance of $\theta_Y^v \cap \theta_Y^w = f_{Y,X}(\theta_X^v \cap \theta_X^w)$. Combining this with the above finding that $\Pi_Y^v(\theta^w \cap \bar{\Sigma}^{\xi^v})$ lies within bounded distance of $\theta_Y^v \cap \theta_Y^w$, we finally conclude that $\text{diam}(\Pi_Y^v(\theta^w))$ is uniformly bounded. \square

3.5. Projections. Here we define projections $\Lambda^v : \bar{\Sigma} \rightarrow \mathbb{P}(\mathcal{K}^v)$ and $\xi^v : \bar{\Sigma} \rightarrow \mathbb{P}(\Xi^v)$. In preparation, we first define $\Pi^v : \bar{\Sigma} \rightarrow \mathbb{P}(\partial\Theta^v)$ by

$$\Pi^v(x) = \bigcup_{Y \in \partial B_v} \Pi_Y^v(x), \quad \text{for any } v \in \mathcal{V} \text{ and } x \in \bar{\Sigma}.$$

In words, $\Pi^v(x)$ consists of the v -windows of x in all fibers over ∂B_v . As before, we extend to arbitrary subsets $U \subset \bar{E}$ by setting $\Pi^v(U) = \Pi^v(U \cap \bar{\Sigma})$.

Now, for each $v \in \mathcal{V}$ our projections are defined as the compositions

$$\Lambda^v = \lambda^v \circ \Pi^v \quad \text{and} \quad \xi^v(x) = i^v \circ \Pi^v.$$

A useful observation is that for any two vertices $v, w \in \mathcal{V}$ with $d_{\text{tree}}(v, w) = 1$ and $X \in \bar{D}$, we have

$$(8) \quad \xi^v(\theta_X^w) = \xi^v(\theta^w) = w \in \Xi^v.$$

Mirroring the notation $d_Y^v(x, y)$ above, for $x, y \in \bar{\Sigma}$ we define

$$d_{\mathcal{K}^v}(x, y) = \text{diam}(\Lambda^v(x) \cup \Lambda^v(y)) \quad \text{and} \quad d_{\Xi^v}(x, y) = \text{diam}(\xi^v(x) \cup \xi^v(y)).$$

Lemma 3.18. *There is a function $N' : [0, \infty) \rightarrow [0, \infty)$ such that for all $v \in \mathcal{V}$, $X \in \partial B_v$, and $x, y \in \bar{\Sigma}$ the quantities $d_{\mathcal{K}^v}(x, y)$ and $d_{\Xi^v}(x, y)$ are at most $N'(d_X^v(x, y))$.*

Proof. For $Y \in \partial B_v$, let us compare the images of some subset $U \subset \partial\Theta_Y^v$ and

$$\text{Smear}(U) = \bigcup_{Z \in \partial B_v} f_{Z,Y}(U) \subset \partial\Theta^v$$

under the maps λ^v, i^v . Since the boundary components of Θ^v are preserved by the maps $f_{Z,Y}$, the images $i^v(U) = i^v(\text{Smear}(U))$ are exactly the same. Moreover, for each $x \in \partial\Theta_Y^v$ we have $\lambda^v(\text{Smear}(\{x\})) = \lambda^v(\ell_{x, \alpha(v)})$, and so by Proposition 3.1(2) this is a set with diameter at most K_1 . Therefore $\lambda^v(U)$ and $\lambda^v(\text{Smear}(U))$ have Hausdorff distance at most $2K_1$.

Now, let $x, y \in \bar{\Sigma}$ and $X \in \partial B_v$ be as in the statement. Set $U = \Pi_X^v(x) \cup \Pi_X^v(y)$ so that $\text{diam}(U) = d_X^v(x, y)$. Since $\Pi^v(x) = \text{Smear}(\Pi_X^v(x))$ and similarly for $\Pi^v(y)$, we see that

$$d_{\mathcal{K}^v}(x, y) = \text{diam}(\lambda^v(\text{Smear}(U))) \quad \text{and} \quad d_{\Xi^v}(x, y) = \text{diam}(i^v(\text{Smear}(U))).$$

Since λ^v is coarsely K_1 -Lipschitz by Proposition 3.11, the preceding paragraph and the triangle inequality shows that

$$d_{\mathcal{K}^v}(x, y) = \text{diam}(\lambda^v(\text{Smear}(U))) \leq \text{diam}(\lambda^v(U)) + 2K_1 \leq (K_1 d_X^v(x, y) + K_1) + 2K_1.$$

Similarly, by Corollary 3.9 we have that

$$\text{diam}_{\Xi^v}(x, y) = \text{diam}(i^v(U)) \leq N(d_X^v(x, y)) + 1.$$

Setting $N'(t) = \max\{N(t) + 1, K_1 t + 3K_1\}$ completes the proof. \square

Proposition 3.19. *There exists $K_4 > 0$ so that for any $v \in \mathcal{V}$:*

- (1) Λ^v and ξ^v are K_4 -coarsely Lipschitz;
- (2) For any $w \in \mathcal{V}$, we have
 - (a) $\text{diam}(\Lambda^v(\Theta^w)) < K_4$, unless $d_{\text{tree}}(v, w) \leq 1$;
 - (b) $\text{diam}(\xi^v(\Theta^w)) < K_4$ unless $w = v$.

Proof. For part (1), we first observe that by [DDLS21, Lemma 3.5], there exists an $R > 0$ so that any pair of points $x, y \in \bar{\Sigma}$ may be connected by a path of length at most $R\bar{d}(x, y)$ that is a concatenation of at most $R\bar{d}(x, y) + 1$ pieces, each of which is either a saddle connection of length at most R in a vertical vertical fiber, or a horizontal geodesic segment in \bar{E} . By the triangle inequality, it thus suffices to assume that x and y are the endpoints of either a horizontal geodesic or a vertical saddle connection of length at most R . Appealing to Lemma 3.18, it further suffices to show that $d_X^v(x, y)$ is linearly bounded by $\bar{d}(x, y)$ for some $X \in \partial B_v$. Lemmas 3.11 and 3.12(1) handle the horizontal segment case, since we are free to subdivide such a path into $\lceil \bar{d}(x, y) \rceil$ segments of length at most 1.

For the vertical segment case we assume x and y lie in the same fiber E_Y and differ by a saddle connection δ of length at most R . Let θ_Y^w , where $w \in \mathcal{V}$, be the spine containing δ . The fact that δ is bounded means that $Y = \pi(x) = \pi(y)$ is bounded distance from the horocycle ∂B_w . Let $X = c_w(Y) \in \partial B_w$ and let $\delta' = f_{X,Y}(\delta)$ be the saddle connection in θ_X^w connecting $x' = f_{X,Y}(x)$ and $y' = f_{X,Y}(y)$. By the triangle inequality and the first part above about bounded length horizontal segments, it suffices to work with the points $x', y' \in \theta^w$. There are three cases to consider: Firstly, if $v = w$, then $x', y' \in \theta_X^v$ so that $\Pi_X^v(x')$ and $\Pi_X^v(y')$ choose the closest cone points in $\partial \Theta_X^v$ to x' and y' , respectively. Since x' and y' are close, so are $\Pi_X^v(x')$ and $\Pi_X^v(y')$. Secondly, if $d_{\text{tree}}(v, w) > 1$, then Corollaries 3.15 and 3.16, and Lemma 3.17 give a uniform bound on $d_Z^v(x', y') \leq \text{diam}(\Pi_Z^v(\theta^w))$ for any point $Z \in c_v(\partial B_w)$. Finally, if $d_{\text{tree}}(v, w) = 1$ then θ_X^w and θ_X^v are adjacent non-crossing spines in E_X . Since θ_X^w is totally geodesic, it follows that $\Pi_X^v(x')$ and $\Pi_X^v(y')$ are either equal or connected by a single edge of $\partial \Theta_X^v$. But this saddle connection has uniformly bounded length, since $X \in \partial B_v$, which completes the proof of (1).

For (2), first recall that strips/saddle connections in the $\alpha(w)$ direction have uniformly bounded width/length over ∂B_w (Lemma 2.3). Therefore $\Theta^w \cap \bar{\Sigma}$ is contained in a bounded neighborhood of $\theta^w \cap \bar{\Sigma}$. By part (1) it thus suffices to bound $\text{diam}(\Lambda^v(\theta^w))$ and $\text{diam}(\xi^v(\theta^w))$. When $d_{\text{tree}}(v, w) \geq 2$, Corollaries 3.15 and 3.16, and Lemma 3.17, imply that there exists $X \in \partial B_v$ so that $\Pi_X^v(\theta^w)$ has bounded diameter in Θ_X^v . Appealing to Lemma 3.18 now bounds $\text{diam}(\Lambda^v(\theta^w))$ and $\text{diam}(\xi^v(\theta^w))$ in these cases. For the remaining case $d_{\text{tree}}(v, w) = 1$ of 2(b), we note that $\xi^v(\theta^w)$ is a single point by (8), and thus 2(b) follows. \square

4. HIERARCHICAL HYPERBOLICITY OF Γ

In this section we complete the proof that Γ is hierarchically hyperbolic. We will use a criterion from [BHMS20], which we now briefly discuss. For further information and heuristic discussion of this approach to hierarchical hyperbolicity, we refer the reader to [BHMS20, §1.5, “User’s guide and a simple example”].

Consider a simplicial complex \mathcal{X} and a graph \mathcal{W} whose vertex set is the set of maximal simplices of \mathcal{X} . The pair $(\mathcal{X}, \mathcal{W})$ is called a *combinatorial HHS* if it satisfies the requirements listed in Definition 4.8 below, and [BHMS20, Theorem 1.18] guarantees that in this case \mathcal{W} is an HHS. The main requirement is along the lines of: \mathcal{X} is hyperbolic, and links of simplices of \mathcal{X} are also hyperbolic. However, this is rarely the case because co-dimension–1 faces of maximal simplices have discrete links. To rectify this, additional edges (coming from \mathcal{W}) should be added to \mathcal{X} and its links as detailed in Definition 4.2. In our case, after adding these edges, \mathcal{X} will be quasi-isometric to \hat{E} , and each other link will be quasi-isometric to either a point or to one of the spaces \mathcal{K}^v or Ξ^v introduced in §3.

There are two natural situations where such pairs arise that the reader might want to keep in mind. First, consider a group H acting on a simplicial complex \mathcal{X} so that there is one orbit of maximal simplices, and those have trivial stabilizers. In this case, we take \mathcal{W} to be (a graph isomorphic to) a Cayley graph of H . (More generally, if the action is cocompact with finite stabilizers of maximal simplices, then the appropriate \mathcal{W} is quasi-isometric to a Cayley graph.) For the second situation, \mathcal{X} is the curve graph of a surface; then maximal simplices are pants decompositions of the surface and \mathcal{W} can be taken to be the pants graph. We will use this as a working example below, when we get into the details.

Most of the work carried out in §3 will be used (as a black-box) to prove that, roughly, links are quasi-isometrically embedded in a space obtained by removing all the “obvious” vertices that provide shortcuts between vertices of the link. This can be seen as an analogue of Bowditch’s fineness condition in the context of relative hyperbolicity.

This section is organized as follows. In §4.1 we list all the relevant definitions and results from [BHMS20], and we illustrate them using pants graphs. In §4.2 we construct the relevant combinatorial HHS for our purposes. In §4.3 we analyze all the various links and related combinatorial objects; we note that most of the work done in §3 is used here to prove Lemma 4.22. At that point, essentially only one property of combinatorial HHSs will be left to be checked, and we do so in §4.4.

4.1. Basic definitions. We start by recalling some basic combinatorial definitions and constructions. Let X be a flag simplicial complex.

Definition 4.1 (Join, link, star). Given disjoint simplices Δ, Δ' of X , the *join* is denoted $\Delta \star \Delta'$ and is the simplex spanned by $\Delta^{(0)} \cup \Delta'^{(0)}$, if it exists. More generally, if K, L are disjoint induced subcomplexes of X such that every vertex of K is adjacent to every vertex of L , then the join $K \star L$ is the induced subcomplex with vertex set $K^{(0)} \cup L^{(0)}$.

For each simplex Δ , the *link* $\text{Lk}(\Delta)$ is the union of all simplices Δ' of X such that $\Delta' \cap \Delta = \emptyset$ and $\Delta' \star \Delta$ is a simplex of X . The *star* of Δ is $\text{Star}(\Delta) = \text{Lk}(\Delta) \star \Delta$, i.e. the union of all simplices of X that contain Δ .

We emphasize that \emptyset is a simplex of X , whose link is all of X and whose star is all of X .

Definition 4.2 (X -graph, W -augmented dual complex). An X -graph is any graph W whose vertex set is the set of maximal simplices of X (those not contained in any larger simplex).

For a flag complex X and an X -graph W , the W -augmented dual graph X^{+W} is the graph defined as follows:

- the 0-skeleton of X^{+W} is $X^{(0)}$;
- if $v, w \in X^{(0)}$ are adjacent in X , then they are adjacent in X^{+W} ;
- if two vertices in W are adjacent, then we consider σ, ρ , the associated maximal simplices of X , and in X^{+W} we connect each vertex of σ to each vertex of ρ .

We equip W with the usual path-metric, in which each edge has unit length, and do the same for X^{+W} . Observe that the 1-skeleton of X is a subgraph $X^{(1)} \subset X^{+W}$.

We provide a running example to illustrate the various definitions in a familiar situation. This example will not be used in the sequel.

Example 4.3. If \mathcal{X} is the curve complex of the surface S , then an example of the an X -graph, W , is the pants graph, since a maximal simplex is precisely a pants decomposition. The W -augmented dual graph can be thought of as adding to the curve graph, $X^{(0)}$ an edge between any two curves that fill a one-holed torus or four-holed sphere and intersect once or twice, respectively: indeed, these subsurfaces are precisely those where an *elementary move* happens as in the definition of adjacency in the pants graph.

Definition 4.4 (Equivalent simplices, saturation). For Δ, Δ' simplices of X , we write $\Delta \sim \Delta'$ to mean $\text{Lk}(\Delta) = \text{Lk}(\Delta')$. We denote by $[\Delta]$ the equivalence class of Δ . Let $\text{Sat}(\Delta)$ denote the set of vertices $v \in X$ for which there exists a simplex Δ' of X such that $v \in \Delta'$ and $\Delta' \sim \Delta$, i.e.

$$\text{Sat}(\Delta) = \left(\bigcup_{\Delta' \in [\Delta]} \Delta' \right)^{(0)}.$$

We denote by \mathfrak{S} the set of \sim -classes of non-maximal simplices in X .

Definition 4.5 (Complement, link subgraph). Let W be an X -graph. For each simplex Δ of X , let Y_Δ be the subgraph of X^{+W} induced by the set of vertices $(X^{+W})^{(0)} - \text{Sat}(\Delta)$.

Let $\mathcal{C}(\Delta)$ be the full subgraph of Y_Δ spanned by $\text{Lk}(\Delta)^{(0)}$. Note that $\mathcal{C}(\Delta) = \mathcal{C}(\Delta')$ whenever $\Delta \sim \Delta'$. (We emphasize that we are taking links in X , not in X^{+W} , and then considering the subgraphs of Y_Δ induced by those links.)

We now pause and continue with the illustrative example.

Example 4.6. Let X and W be as in Example 4.3. A simplex Δ is a multicurve which determines two (open) subsurfaces $U = U(\Delta), U' = U'(\Delta) \subset S$, where U is the union of the complementary components of the multicurve that are not a pair of pants, and $U' = S - \bar{U}$. Note that $\partial U \subset \Delta$ is a submulticurve and that $\Delta - \partial U$ is a pants decomposition of U' . A simplex Δ' is equivalent to Δ if it defines the same subsurfaces. Thus $\text{Sat}(\Delta)$ consists of $\partial U(\Delta)$ together with all essential curves in $U'(\Delta)$, while $\mathcal{C}(\Delta)$ is the join of graphs quasi-isometric to curve graphs of the components of $U(\Delta)$. For components of $U(\Delta)$ which are one-holed tori or

four-holed spheres, the corresponding subgraphs are isometric to their curve graphs (since the extra edges in X^{+W} precisely give edges for these curve graphs).

Definition 4.7 (Nesting). Let X be a simplicial complex. Let Δ, Δ' be non-maximal simplices of X . Then we write $[\Delta] \sqsubseteq [\Delta']$ if $\text{Lk}(\Delta) \subseteq \text{Lk}(\Delta')$.

We note that if $\Delta' \subset \Delta$, then $[\Delta] \sqsubseteq [\Delta']$. Also, for Example 4.3, 4.6, $[\Delta] \sqsubseteq [\Delta']$ if and only if $U(\Delta) \subset U(\Delta')$.

Finally, we are ready for the main definition:

Definition 4.8 (Combinatorial HHS). A *combinatorial HHS* (X, W) consists of a flag simplicial complex X and an X -graph W satisfying the following conditions for some $n \in \mathbb{N}$ and $\delta \geq 1$:

- (1) any chain $[\Delta_1] \sqsubset [\Delta_2] \sqsubset \dots$ has length at most n ;
- (2) for each non-maximal simplex Δ , the subgraph $\mathcal{C}(\Delta)$ is δ -hyperbolic and (δ, δ) -quasi-isometrically embedded in Y_Δ ;
- (3) Whenever Δ and Δ' are non-maximal simplices for which there exists a non-maximal simplex Γ such that $[\Gamma] \sqsubseteq [\Delta]$, $[\Gamma] \sqsubseteq [\Delta']$, and $\text{diam}(\mathcal{C}(\Gamma)) \geq \delta$, then there exists a simplex Π in the link of Δ' such that $[\Delta' \star \Pi] \sqsubseteq [\Delta]$ and all $[\Gamma]$ as above satisfy $[\Gamma] \sqsubseteq [\Delta' \star \Pi]$;
- (4) if v, w are distinct non-adjacent vertices of $\text{Lk}(\Delta)$, for some simplex Δ of X , contained in W -adjacent maximal simplices, then they are contained in W -adjacent simplices of the form $\Delta \star \Delta'$.

We will see below that combinatorial HHS give HHSs. The reader not interested in the explicit description of the HHS structure can skip the following two definitions.

Definition 4.9 (Orthogonality, transversality). Let X be a simplicial complex. Let Δ, Δ' be non-maximal simplices of X . Then we write $[\Delta] \perp [\Delta']$ if $\text{Lk}(\Delta') \subseteq \text{Lk}(\text{Lk}(\Delta))$. If $[\Delta]$ and $[\Delta']$ are not \perp -related or \sqsubseteq -related, we write $[\Delta] \pitchfork [\Delta']$.

Definition 4.10 (Projections). Let (X, W) be a combinatorial HHS.

Fix $[\Delta] \in \mathfrak{S}$ and define a map $\pi_{[\Delta]}: W \rightarrow \mathbb{P}(\mathcal{C}([\Delta]))$ as follows. First let $p: Y_\Delta \rightarrow \mathbb{P}(\mathcal{C}([\Delta]))$ be the coarse closest-point projection, i.e.

$$p(x) = \{y \in \mathcal{C}([\Delta]) \mid d_{Y_\Delta}(x, y) \leq d_{Y_\Delta}(x, \mathcal{C}([\Delta])) + 1\}.$$

Suppose that $w \in W^{(0)}$, so w corresponds to a unique simplex Δ_w of X . Define

$$\pi_{[\Delta]}(w) = p(\Delta_w \cap Y_\Delta).$$

We have thus defined $\pi_{[\Delta]}: W^{(0)} \rightarrow \mathbb{P}(\mathcal{C}([\Delta]))$. If $v, w \in W^{(0)}$ are joined by an edge e of W , then Δ_v, Δ_w are joined by edges in X^{+W} , and we let $\pi_{[\Delta]}(e) = \pi_{[\Delta]}(v) \cup \pi_{[\Delta]}(w)$.

Now let $[\Delta], [\Delta'] \in \mathfrak{S}$ satisfy $[\Delta] \pitchfork [\Delta']$ or $[\Delta'] \sqsubset [\Delta]$. Let

$$\rho_{[\Delta]}^{[\Delta']} = p(\text{Sat}(\Delta') \cap Y_\Delta).$$

Let $[\Delta] \sqsubset [\Delta']$. Let $\rho_{[\Delta]}^{[\Delta']}: \mathcal{C}([\Delta']) \rightarrow \mathcal{C}([\Delta])$ be the restriction of p to $\mathcal{C}([\Delta']) \cap Y_\Delta$, and \emptyset otherwise.

The next theorem from [BHMS20] provides the criteria we will use to prove that Γ is a hierarchically hyperbolic group.

Given a combinatorial HHS (X, W) , we denote \mathfrak{S}_W the set as in Definition 4.4, endowed with nesting and orthogonality relations as in Definitions 4.7 and 4.9. Also, we associated to \mathfrak{S}_W the hyperbolic spaces as in Definition 4.8, and define projections as in Definition 4.10.

Theorem 4.11. [BHMS20, Theorem 1.18, Remark 1.19] *Let (X, W) be a combinatorial HHS. Then (W, \mathfrak{S}_W) is a hierarchically hyperbolic space.*

Moreover, if a group G acts by simplicial automorphisms on X with finitely many orbits of links of simplices, and the resulting G -action on maximal simplices extends to a metrically proper cobounded action on W , then G acts metrically properly and coboundedly by HHS automorphisms on (W, \mathfrak{S}_W) , and is therefore a hierarchically hyperbolic group.

4.2. Combinatorial HHS structure. We now define a flag simplicial complex \mathcal{X} . The vertex set is $\mathcal{X}^{(0)} = \mathcal{V} \sqcup \mathcal{K}$, where

$$\mathcal{K} = \bigsqcup_{v \in \mathcal{V}} \mathcal{K}^v.$$

Given a vertex $s \in \mathcal{K}$, let $v(s) \in \mathcal{V}$ be the unique vertex with $s \in \mathcal{K}^v$. We also write $\alpha(s) = \alpha(v(s))$.

There are 3 types of edges (see Figure 5):

- (1) $v, w \in \mathcal{V}$ are connected by an edge if and only if $d_{\text{tree}}(v, w) = 1$.
- (2) $s, t \in \mathcal{K}$ are connected by an edge if and only if $d_{\text{tree}}(v(s), v(t)) = 1$.
- (3) $s \in \mathcal{K}$ and $w \in \mathcal{V}$ are connected by an edge if and only if $d_{\text{tree}}(v(s), w) \leq 1$.

We declare \mathcal{X} to be the flag simplicial complex with the 1-skeleton defined above.




FIGURE 5. The simplicial map Z restricted to a part of \mathcal{X} (on the left) to a part of the union of trees T_α (on the right).

The map $\mathcal{K} \rightarrow \mathcal{V}$ given by $s \mapsto v(s)$ and the identity $\mathcal{V} \rightarrow \mathcal{V}$ extends to a surjective simplicial map

$$Z: \mathcal{X} \rightarrow \bigsqcup_{\alpha \in \mathcal{P}} T_\alpha.$$

We note that we may view the union $\bigsqcup T_\alpha$ on the right as a subgraph of $\mathcal{X}^{(1)}$, making Z a retraction.

For any vertex v in any tree T_α , $Z^{-1}(v)$ is the join of $\{v\}$ and the set \mathcal{K}^v :

$$(9) \quad Z^{-1}(v) = \{v\} \star \mathcal{K}^v.$$

For any pair of adjacent vertices $v, w \in T_\alpha$ (so $d_{\text{tree}}(v, w) = 1$), the preimage of the edge $[v, w] \subset T_\alpha$ is also a join:

$$(10) \quad Z^{-1}([v, w]) = Z^{-1}(v) \star Z^{-1}(w) = (\{v\} \star \mathcal{K}^v) \star (\{w\} \star \mathcal{K}^w).$$

Lemma 4.12. *The maximal simplices of \mathcal{X} are exactly the 3-simplices with vertex set $\{s, v(s), t, v(t)\}$ where $s, t \in \mathcal{K}$ and $d_{\text{tree}}(v(s), v(t)) = 1$. In this case, we say that (s, t) defines a maximal simplex, denoted $\sigma(s, t)$.*

Proof. Because the map Z is simplicial, any simplex of \mathcal{X} is contained in $Z^{-1}(v)$ or $Z^{-1}([v, w])$ for some vertex v in some T_α or some edge $[v, w]$ in some T_α . The lemma thus follows from (9) and (10). \square

Given a vertex $s \in \mathcal{K}$, recall from Lemma 3.5 that $M(s) = (\lambda^{v(s)})^{-1}(N_{K_1}(s)) \subset \Theta^{v(s)}$, for K_1 as in Proposition 3.1 (and Lemma 3.5). Given a pair of vertices (s, t) in \mathcal{K} that define a maximal simplex $\sigma(s, t)$, we will write $M(s, t) = M(s) \cap M(t)$.

Lemma 4.13. *There exists $R > 0$ with the following properties.*

- (1) *For any pair of adjacent vertices $s, t \in \mathcal{K}$ (i.e., defining a maximal simplex $\sigma(s, t)$), $M(s, t)$ is a non-empty subset of diameter at most R .*
- (2) *Given $v \in \mathcal{V}$, we have*

$$\Theta^v = \bigcup_{v(s)=v, d_{\text{tree}}(v(s), v(t))=1} M(s, t).$$

- (3) *Fixing s , we have*

$$M(s) = \bigcup_{d_{\text{tree}}(v(s), v(t))=1} M(s, t).$$

- (4) *The collection of all $M(s, t)$ is R -dense in \bar{E} .*

Proof. Item (1) follows from Proposition 3.1(3). More precisely, the fact that $M(s, t)$ is non-empty follows from K_1 -coarse-surjectivity of $\lambda^{v(s)} \times \lambda^{v(t)}$, while boundedness follows from the fact that said map is a quasi-isometry.

In order to show item (2), notice that

$$\Theta^v = \bigcup_{d_{\text{tree}}(v, w)=1} \Theta^v \cap \Theta^w.$$

That is, every point of Θ^v is also in Θ^w for some w adjacent to v . In view of this, we conclude by noticing that if $x \in \Theta^v \cap \Theta^w$, then $x \in M(\lambda^v(x), \lambda^w(x))$. Item (3) follows similarly.

Finally, item (4) follows from item (2) and the fact that the collection of all Θ^v is coarsely dense in \bar{E} . \square

Next we define a graph \mathcal{W} whose vertex set is the set of maximal simplices of \mathcal{X} . We would like to just connect maximal simplices when the corresponding subsets $M(s, t)$ are close in \bar{E} (first bullet below); however, in order to arrange item (4) of the definition of combinatorial HHS (and only for that reason) we need different closeness constants for different situations. We fix R as in Lemma 4.13, and moreover we require $R > K_1^2 + K_1$, for K_1 as in Proposition 3.1 and Lemma 3.5.

Given maximal simplices $\sigma(s_1, t_1)$ and $\sigma(s_2, t_2)$, we declare them to be connected by an edge in \mathcal{W} if one of the following holds:

- (11)
 - $\bar{d}(M(s_1, t_1), M(s_2, t_2)) \leq 10R$
 - $s_1 = s_2$ and $\bar{d}(M(t_1), M(t_2)) \leq 10R$

Here the \bar{d} -distances are the *infimal* distances between the sets in \bar{E} (as opposed to the diameter of the union). Note that since $M(s, t) = M(t, s)$, the second case also implicitly describes a “symmetric case” with s_i and t_i interchanged.

The following is immediate from Lemma 3.6, setting $R' = \max\{10R, N_{\mathcal{K}}(10R)\}$.

Lemma 4.14. *There exists $R' \geq 10R$ so then the following holds. If $s, t_1, t_2 \in \mathcal{K}$ are vertices with s connected to both t_i in \mathcal{X} and $\bar{d}(M(t_1), M(t_2)) \leq 10R$, then $\bar{d}(M(s, t_1), M(s, t_2)) \leq R'$. In particular, whenever $\sigma(s_1, t_1)$ and $\sigma(s_2, t_2)$ are connected in \mathcal{W} , we have $\bar{d}(M(s_1, t_1), M(s_2, t_2)) \leq R'$.*

Lemma 4.15. *\mathcal{W} is quasi-isometric to \bar{E} , by mapping each $\sigma(s, t)$ to (any point in) $M(s, t)$. Moreover, the extension group Γ acts by simplicial automorphisms on \mathcal{X} , induced by the existing action on $\mathcal{V} \subseteq \mathcal{X}^{(0)}$ and the action on $\mathcal{K} \subseteq \mathcal{X}^{(0)}$ as in Proposition 3.1(4). The resulting action on maximal simplices extends to a metrically proper cobounded action on \mathcal{W} .*

Proof. In view of Lemma 4.14, the first part follows by combining Lemma 4.13(4) and Proposition 2.4 (applied to any choice of a point in each $M(s, t)$).

It is immediate to check that the Γ -action defined on the 0-skeleton of \mathcal{X} extends to an action on \mathcal{X} . That the resulting action on maximal simplices of \mathcal{X} (that is, the 0-skeleton of \mathcal{W}) extends to an action on \mathcal{W} follows from the equivariance property in Proposition 3.1(4) and the definitions of the sets $M(s)$ and $M(s, t)$.

Moreover, the quasi-isometry $\mathcal{W} \rightarrow \bar{E}$ described in the statement is Γ -equivariant, so that the action of Γ on \mathcal{W} is metrically proper and cobounded since the action of Γ on \bar{E} has these properties. \square

The goal for the remainder of this section is to prove the following.

Theorem 4.16. *The pair $(\mathcal{X}, \mathcal{W})$ is a combinatorial HHS. Moreover, there is an action of Γ on \mathcal{X} satisfying the properties stated in Theorem 4.11.*

4.3. Simplices, links, and saturations. Before giving the proof of Theorem 4.16, we begin by describing explicitly the kinds of simplices of \mathcal{X} that there are, explain what their links and saturations are, and observe some useful properties.

Lemma 4.17 (Empty simplex). *For the empty simplex, $\mathcal{C}(\emptyset) = \mathcal{X}^{+\mathcal{W}}$ is quasi-isometric to \hat{E} .*

Proof. We define a map $Z': \mathcal{X}^{+\mathcal{W}} \rightarrow \hat{E}$ that extends the (restricted) simplicial map $Z: \mathcal{X}^{(1)} \rightarrow \bigsqcup T_\alpha$ already constructed above. To do that, we must extend over each edge $e = [x, y]$ of $\mathcal{X}^{+\mathcal{W}}$ coming from the edge of \mathcal{W} connecting $\sigma(s_1, t_1)$ and $\sigma(s_2, t_2)$. Since $Z(x), Z(y) \in \mathcal{V}$, and $\bar{d}(M(s_1, t_1), M(s_2, t_2)) \leq R'$ (for R' as in Lemma 4.14), we see that $d_{\hat{E}}(v, w) \leq R'$. We can then define Z' on e to be a constant speed parameterization of a uniformly bounded length path from $Z(x)$ to $Z(y)$. It follows that Z' is Lipschitz.

The union of the trees $\bigsqcup T_\alpha$ is R_0 -dense for some $R_0 > 0$ by [DDLS21, Lemma 3.6], so it suffices to find a one-sided inverse to Z' , from $\bigsqcup T_\alpha$ to $\mathcal{X}^{+\mathcal{W}}$, and show that with respect to the subspace metric from \hat{E} , it is coarsely Lipschitz. As already noted, Z restricts to a retraction of $\mathcal{X}^{(1)}$ onto $\bigsqcup T_\alpha \subset \mathcal{X}^{(1)} \subset \mathcal{X}^{+\mathcal{W}}$, which is thus the required one-sided inverse. All that remains is to show that it is coarsely Lipschitz.

According to [DDLS21, Lemma 3.8], any $v \in T_\alpha, w \in T_\beta$ are connected by a *combinatorial path* of length comparable to $\hat{d}(v, w)$. Such a path is the concatenation

of *horizontal jumps*, each of which is the \bar{P} -image in \hat{E} of a geodesic in \bar{D}_z , for some $z \in \bar{\Sigma}$, that connects two components of $\partial\bar{D}_z$ and whose interior is disjoint from $\partial\bar{D}_z$. From that same lemma, we may assume each horizontal jump has length uniformly bounded above and below, and thus has total number of jumps bounded in terms of $\hat{d}(v, w)$. Therefore, we can reduce to the case that v, w are joined by a single horizontal jump of bounded length. Such a horizontal jump can also be regarded as a path in \bar{E} connecting Θ^v to Θ^w . Hence, in view of Lemma 4.13(2), there are $M(s_1, t_1) \subseteq \Theta^v$ and $M(s_2, t_2) \subseteq \Theta^w$ within uniformly bounded distance of each other in \bar{E} . Lemma 4.15 implies that there exists a path in \mathcal{W} of uniformly bounded length from $\sigma(s_1, t_1)$ to $\sigma(s_2, t_2)$, which can be easily turned into a path of uniformly bounded length from v to w in $\mathcal{X}^{+\mathcal{W}}$, as required. \square

There is an important type of 1-dimensional simplex, which we call a Ξ -type simplex, due to the following lemma. See Figure 6. Given $w \in \mathcal{V}$, set

$$(12) \quad \text{Lk}_{\Xi}(w) = \bigcup_{d_{\text{tree}}(u, w)=1} Z^{-1}(u) = \bigcup_{d_{\text{tree}}(u, w)=1} \{u\} \star \mathcal{K}^u.$$

Lemma 4.18 (Ξ -type simplex). *Let Δ be a 1-simplex of \mathcal{X} with vertices $s, v(s)$, for $s \in \mathcal{K}$. Then*

$$\text{Lk}(\Delta) = \text{Lk}_{\Xi}(v(s)) \quad \text{and} \quad \text{Sat}(\Delta) = \{v(s)\} \cup \mathcal{K}^{v(s)}.$$

Moreover, $\mathcal{C}(\Delta)$ is quasi-isometric to $\Xi^{v(s)}$, via a quasi-isometry which is the identity on $\mathcal{V} \cap \mathcal{C}(\Delta)$ and maps t to $v(t)$ for $t \in \mathcal{K} \cap \mathcal{C}(\Delta)$.




FIGURE 6. Ξ -type simplex: Part of the link and saturation of a 1-simplex Δ with vertices $s, v = v(s)$.

Proof. It is clear from the definitions that the link of Δ is as described. Also, any simplex with vertex set of the form $\{v(s), t\}$ for some $t \in \mathcal{K}^{v(s)}$ has the same link as Δ . Therefore, to prove that the saturation is as described we are left to show if a simplex Δ' has the same link as Δ , then its vertex set is contained in the set we described. If $w \in \mathcal{V}$ is a vertex of Δ' , then $d_{\text{tree}}(v, w) = 1$ for all neighbors v of $v(s)$ in $T_{\alpha(v)}$. This implies $w = v(s)$. Similarly, if $t \in \mathcal{K}$ is a vertex of Δ' , then $v(t) = v(s)$, and we are done.

Let us show that the map given in the statement is coarsely Lipschitz. To do so, it suffices to consider $w \neq w'$ with $d_{\text{tree}}(w, v(s)) = d_{\text{tree}}(w', v(s)) = 1$ and connected by an edge in $\mathcal{X}^{+\mathcal{W}}$, and show that they are connected by a bounded-length path in Ξ^v . We argue below that $\bar{d}(\Theta^w \cap \Theta^{v(s)}, \Theta^{w'} \cap \Theta^{v(s)}) \leq R'$. Once we do that, the existence of the required bounded-length path follows directly from Lemma 3.8.

Let us now prove the desired inequality. Notice that w, w' cannot be connected by an edge of \mathcal{X} since they are both distance 1 from $v(s)$ in the tree $T_{\alpha(v)}$. Hence, w and w' are contained respectively in maximal simplices $\sigma(t, u)$ and $\sigma(t', u')$ connected by an edge in \mathcal{W} . Say, up to swapping t with u and/or t' with u' , that $v(t) = w$ and $v(t') = w'$. In either case of the definition of the edges of \mathcal{W} we have $\bar{d}(M(t), M(t')) \leq 10R$. Since $M(t, s) \subseteq \Theta^w \cap \Theta^{v(s)}$ and $M(t', s) \subseteq \Theta^{w'} \cap \Theta^{v(s)}$, using Lemma 4.14 we get

$$\bar{d}(\Theta^w \cap \Theta^{v(s)}, \Theta^{w'} \cap \Theta^{v(s)}) \leq \bar{d}(M(t, s), M(t', s)) \leq R',$$

as we wanted.

Conversely, if w, w' as above are joined by an edge in $\Xi^{v(s)}$ we will now show that they are also connected by an edge in $\mathcal{C}(\Delta)$. By definition of $\Xi^{v(s)}$, we have $\Theta^w \cap \Theta^{w'} \neq \emptyset$. By Lemma 4.13(2), There exists t, u, t', u' with $v(t) = w$ and $v(t') = w'$, so that $M(t, u) \cap M(t', u') \neq \emptyset$. In particular, $\bar{d}(M(t, u), M(t', u')) \leq 10R$. This says that $\sigma(t, u)$ and $\sigma(t', u')$ are connected in \mathcal{W} , and hence that w and w' are connected in $\mathcal{X}^{+\mathcal{W}}$, as required. \square

There is also an important type of 2-dimensional simplex, which we call a \mathcal{K} -type simplex, due to the next lemma. See Figure 7.

Lemma 4.19 (\mathcal{K} -type simplex). *Let Δ be a 2-simplex of \mathcal{X} with vertices $s, v(s), w$, for $s \in \mathcal{K}$ and $w \in \mathcal{V}$ with $d_{\text{tree}}(w, v(s)) = 1$. Then*

$$\text{Lk}(\Delta) = \mathcal{K}^w \quad \text{and} \quad \text{Sat}(\Delta) = \{w\} \cup \text{Lk}_{\Xi}(w)^{(0)}.$$

Moreover, $\mathcal{C}(\Delta)$ is quasi-isometric to \mathcal{K}^w , the quasi-isometry being the identity at the level of vertices.




FIGURE 7. \mathcal{K} -type simplex: Part of the link and saturation of a 2-simplex Δ with vertices $s, v = v(s), w$.

Proof. It is clear from the definitions that the link of Δ is as described. Also, any simplex with vertex set of the form $\{t, v(t), w\}$ for some $t \in \mathcal{K}$ with $d_{\text{tree}}(v(t), w) = 1$ has the same link as Δ . Therefore, to prove that the saturation is as described we are left to show if a simplex Δ' has the same link as Δ , then its vertex set is contained in the set we described. Given any vertex $u \in \mathcal{V}$ of Δ' , it has to be connected to all $t \in \mathcal{K}$ with $v(t) = w$, implying that either $u = w$ or $d_{\text{tree}}(u, w) = 1$, as required for vertices in \mathcal{V} . Similarly, any vertex $u \in \mathcal{K}$ of Δ' has to be connected to all $t \in \mathcal{K}$ with $v(t) = w$, implying $d_{\text{tree}}(v(u), w) = 1$, and we are done.

To prove that $\mathcal{C}(\Delta)$ is naturally quasi-isometric to \mathcal{K}^w , it suffices to show that if $u, t \in \mathcal{C}(\Delta)$ are connected by an edge in $\mathcal{X}^{+\mathcal{W}}$, then they are uniformly close in \mathcal{K}^w and that, vice versa, if $d_{\mathcal{K}^w}(u, t) = 1$, then u, t are connected by an edge in $\mathcal{C}(\Delta)$.

First, if $u, t \in \mathcal{C}(\Delta)$ are connected by an edge in $\mathcal{X}^{+\mathcal{W}}$, then there exist u', t' so that $\bar{d}(M(u, t'), M(u', t)) \leq R'$ (see Lemma 4.14). In particular, $\bar{d}(M(u), M(t)) \leq R'$, which in turn gives a uniform bound on the distance in the path metric of Θ^w between $M(u)$ and $M(t)$ because the metrics \bar{d} and the path metric on θ^w are coarsely equivalent. By Proposition 3.1(1) we must also have a uniform bound on $d_{\mathcal{K}^w}(u, t)$.

Now suppose that $d_{\mathcal{K}^w}(u, t) = 1$. We can then deduce from Proposition 3.1(3) that $\bar{d}(M(s', u), M(s', t)) \leq K_1^2 + K_1 \leq 10R$ for any fixed $s' \in \mathcal{K}^{v(s)}$ (the proposition yields the analogous upper bound in the path metric of Θ^w , which is a stronger statement). Therefore u, t are connected by an edge in $\mathcal{X}^{+\mathcal{W}}$, whence in $\mathcal{C}(\Delta)$, as required. \square

The remaining simplices are not particularly interesting as their links are joins (or points), and hence have diameter at most 3, but we will still need to verify properties for them. We describe each type and its link in the next lemma. Recall the definition of $\text{Lk}_{\Xi}(w)$ in (12).

Lemma 4.20. *The following is a list of all types of non-empty, nonmaximal, simplices Δ of \mathcal{X} that are not of Ξ -type or \mathcal{K} -type, together with their links. In each case, the link is a nontrivial join (or a point), and $\mathcal{C}(\Delta)$ has diameter at most 3.*

In the table below the simplices Δ have vertices $u, w \in \mathcal{V}$ with $d_{\text{tree}}(u, w) = 1$ and $s, t \in \mathcal{K}$ with $d_{\text{tree}}(v(s), v(t)) = 1$ and $d_{\text{tree}}(v(s), u) = 1$.

Δ	$\text{Lk}(\Delta)$
$\{u\}$	$\mathcal{K}^u \star \text{Lk}_{\Xi}(u)$
$\{s\}$	$\{v(s)\} \star \text{Lk}_{\Xi}(v(s))$
$\{u, w\}$	$\mathcal{K}^u \star \mathcal{K}^w$
$\{s, t\}$	$\{v(s), v(t)\}$
$\{s, u\}$	$\{v(s)\} \star \mathcal{K}^u$
$\{s, t, v(s)\}$	$\{v(t)\}$

Proof. This is straightforward given the definition of \mathcal{X} and we leave its verification to the reader. Referring to Figure 5, and comparing with Figures 6 and 7, may be helpful. \square

The next lemma collects a few additional properties we will need. By the *type* of a simplex in \mathcal{X} , we mean its orbit by the action of the simplicial automorphism group of \mathcal{X} . There are 9 types of nonempty simplices: maximal, Ξ -type, \mathcal{K} -type, and the six types listed in Lemma 4.20.

Lemma 4.21. *The following hold in \mathcal{X} .*

- (a) *The link of a simplex with a given type cannot be strictly contained in the link of a simplex with the same type.*
- (b) *For all non-maximal simplices Δ and Δ' so that there is a simplex Γ with $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Delta') \cap \text{Lk}(\Delta)$ and $\text{diam}(\mathcal{C}(\Gamma)) > 3$, there exists a simplex Π in the link of Δ' with $\text{Lk}(\Delta' * \Pi) \subseteq \text{Lk}(\Delta)$ so that for any Γ as above we have $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Delta' * \Pi)$.*

Proof. Part (a) follows directly from the descriptions of the simplices given in Lemmas 4.18, 4.19, and 4.20, and we leave this to the reader.

Before we prove (b), we suppose $\text{Lk}(\Delta') \cap \text{Lk}(\Delta) \neq \emptyset$, and make a few observations. First, Δ' and Δ must project by Z to the same tree: $Z(\Delta'), Z(\Delta) \subset T_\alpha$ for some $\alpha \in \mathcal{P}$. Next, note that $Z(\text{Lk}(\Delta') \cap \text{Lk}(\Delta))$ is contained in the intersection of the stars in T_α of $Z(\Delta')$ and $Z(\Delta)$. Moreover, (as in any tree) the intersection of these two stars is contained in a single edge, or else $Z(\Delta') = Z(\Delta) = \{w\} \in T_\alpha^{(0)} \subset \mathcal{V}$ is a single point. In this latter case, by (9), we have

$$\Delta, \Delta' \subset \{w\} \star \mathcal{K}^w.$$

Next, note that for any \mathcal{K} -type simplex $\Gamma = \{s, v(s), w\}$, $\text{Lk}(\Gamma) = \mathcal{K}^w$ by Lemma 4.19, and if $\text{Lk}(\Gamma) \subset \text{Lk}(\Delta') \cap \text{Lk}(\Delta)$, then w is in the intersection of the stars of $Z(\Delta')$ and $Z(\Delta)$. For a Ξ -type simplex $\Gamma = \{s, v(s)\}$, $\text{Lk}(\Gamma) = \text{Lk}_\Xi(v(s))$ by Lemma 4.18, and together with Lemma 4.20 and the previous paragraph, we see that $\text{Lk}(\Gamma) \subset \text{Lk}(\Delta') \cap \text{Lk}(\Delta)$ if and only if $Z(\Delta') = Z(\Delta) = \{v(s)\}$ in T_α .

With these observations in hand, we proceed to the proof of (b), which divides into two cases.

Case 1. There is a Ξ -type simplex $\Gamma = \{s, v(s)\}$ with $\text{Lk}(\Gamma) \subset \text{Lk}(\Delta') \cap \text{Lk}(\Delta)$.

In this situation, $Z(\Delta') = Z(\Delta) = \{v(s)\}$, and thus $\Delta, \Delta' \subset \{v(s)\} \star \mathcal{K}^{v(s)}$. From Lemmas 4.18 and 4.20, we see that $\text{Lk}(\Delta') \cap \text{Lk}(\Delta)$ must be equal to one of $\text{Lk}(\Gamma)$, $\text{Lk}(\{s\})$, or $\text{Lk}(\{v(s)\})$. Inspection of these links shows that $\text{Lk}(\Gamma)$ is the only link of a Ξ -type simplex contained in it. First suppose that $\text{Lk}(\Delta') \cap \text{Lk}(\Delta)$ has the form $\text{Lk}(\Gamma)$ or $\text{Lk}(\{s\})$. In this situation, we easily find $\Pi \subset \text{Lk}(\Delta')$ so that $\text{Lk}(\Gamma) = \text{Lk}(\Delta' * \Pi)$. Furthermore, for any \mathcal{K} -type simplex link \mathcal{K}^w in the intersection, we must have $\mathcal{K}^w \subset \text{Lk}(\Gamma) = \text{Lk}(\Delta' * \Pi)$. Therefore, the link of any Ξ - or \mathcal{K} -type simplex contained in $\text{Lk}(\Delta') \cap \text{Lk}(\Delta)$ must be contained in $\text{Lk}(\Gamma) = \text{Lk}(\Delta' * \Pi)$, as required. Now suppose instead that $\text{Lk}(\Delta') \cap \text{Lk}(\Delta) = \text{Lk}(\{v(s)\})$. By Lemma 4.20, we see that $\Delta = \Delta' = \{v(s)\}$. In this case, setting $\Pi = \emptyset$ trivially completes the proof since then $\text{Lk}(\Delta') \cap \text{Lk}(\Delta) = \text{Lk}(\Delta') = \text{Lk}(\Delta' * \Pi)$.

Case 2. No link of a Ξ -type simplex is contained in $\text{Lk}(\Delta') \cap \text{Lk}(\Delta)$.

From the observations above, $Z(\Delta)$ and $Z(\Delta')$ do not consist of the same single point, and hence the stars of $Z(\Delta)$ and $Z(\Delta')$ intersect in either a point or an edge in T_α . Since $Z(\text{Lk}(\Delta') \cap \text{Lk}(\Delta))$ is contained in the intersection of these stars, there are at most two \mathcal{K} -type simplices whose links are contained in $\text{Lk}(\Delta') \cap \text{Lk}(\Delta)$. If there are two \mathcal{K} -type simplices Γ, Γ' with $\mathcal{K}^w = \text{Lk}(\Gamma)$ and $\mathcal{K}^u = \text{Lk}(\Gamma')$ contained

in $\text{Lk}(\Delta') \cap \text{Lk}(\Delta)$, then observe that

$$\mathcal{K}^u, \mathcal{K}^w \subset \text{Lk}(\{u, w\}) = \mathcal{K}^u \star \mathcal{K}^w \subset \text{Lk}(\Delta') \cap \text{Lk}(\Delta).$$

By inspection of the possible links in Lemmas 4.18, 4.19, and 4.20, it must be that Δ' is either $\{u\}$, $\{w\}$, or $\{u, w\}$, and so setting Π to be $\{w\}$, $\{u\}$, or \emptyset , respectively, we are done. On the other hand, if there is exactly one \mathcal{K} -type simplex Γ with $\mathcal{K}^w = \text{Lk}(\Gamma) \subset \text{Lk}(\Delta') \cap \text{Lk}(\Delta)$, then again inspecting all possible situations, we can find $\Pi \subset \text{Lk}(\Delta')$ with $\mathcal{K}^w = \text{Lk}(\Gamma) = \text{Lk}(\Delta' \star \Pi)$, and again we are done with this case. This completes the proof. \square

Lemma 4.22. *There exists $L \geq 1$ so that for every non-maximal simplex, there is an (L, L) -coarsely Lipschitz retraction $r_\Delta: Y_\Delta \rightarrow \mathbb{P}(\mathcal{C}(\Delta))$. In particular, $\mathcal{C}(\Delta)$ is uniformly quasi-isometrically embedded in Y_Δ .*

Proof. By Lemma 4.20, we only have to consider simplices of Ξ - and \mathcal{K} -type.

Consider $\Delta = \{s, v\}$ with $v = v(s)$ of Ξ -type first. Recall from Lemma 4.18 that Ξ^v naturally includes into $\mathcal{C}(\Delta)$ by a quasi-isometry. Here we will use make use of the map ξ^v , whose relevant properties for our current purpose are stated in Proposition 3.19. For a vertex $u \in \mathcal{V} \setminus \text{Sat}(\Delta)$ (so, $u \neq v$) we define $r_\Delta(u) = \xi^v(\Theta^u)$. For $t \in \mathcal{K} \setminus \text{Sat}(\Delta)$ we define $r_\Delta(t) = \xi^v(\Theta^{v(t)})$. Notice that the sets $r_\Delta(u)$ are uniformly bounded by Proposition 3.19 (and Lemma 4.18). Also, r_Δ is coarsely the identity on the vertices of $\text{Lk}(\Delta)$ in \mathcal{V} by Equation (8) and Proposition 3.19(2b). To check that r_Δ is coarsely Lipschitz it suffices to consider $\mathcal{X}^{+\mathcal{W}}$ -adjacent vertices of \mathcal{V} . Notice that vertices $w, w' \in \mathcal{V}$ that are adjacent in $\mathcal{X}^{+\mathcal{W}}$ have corresponding $\Theta^w, \Theta^{w'}$ within $10R$ of each other in \bar{E} . Indeed, Θ^v and Θ^w actually intersect if w, w' are adjacent in \mathcal{X} , and they contain subsets $M(\cdot)$ within $10R$ of each other if w, w' are contained in \mathcal{W} -adjacent maximal simplices (this is true regardless of which case of the definition (11) for the edges of \mathcal{W} applies). The fact that r_Δ is coarsely Lipschitz now follows from, Proposition 3.19(1), which says that ξ^v is coarsely Lipschitz on \bar{E} .

Next, consider $\Delta = \{s, v(s), w\}$ of \mathcal{K} -type. For a vertex $u \in \mathcal{V} \setminus \text{Sat}(\Delta)$ (so, $d_{\text{tree}}(u, w) \geq 2$), define $r_\Delta(u) = \Lambda^w(\Theta^u)$. For a vertex $t \in \mathcal{K} \setminus \text{Sat}(\Delta)$, define $r_\Delta(t) = \Lambda^w(M(t))$. Notice that, by definition of $M(t)$, if $t \in \mathcal{K}^w$, then $r_\Delta(t)$ lies within Hausdorff distance K_1 of t . Also, since $d_{\text{tree}}(u, w) \geq 2$, Proposition 3.19(2a) ensures that the diameter of $\Lambda^w(\Theta^u)$ is bounded. Since $M(t) \subseteq \Theta^{v(t)}$, we see that all the sets in the image of r_Δ are bounded, and also we see that in order to prove that r_Δ is Lipschitz it suffices to consider vertices of \mathcal{K} . But vertices $s, t \in \mathcal{K}$ that are adjacent in $\mathcal{X}^{+\mathcal{W}}$ have corresponding $M(s), M(t)$ within $10R$ of each other in \bar{E} , so the conclusion follows from Proposition 3.19(1), which states that Λ^w is coarsely Lipschitz on \bar{E} . \square

4.4. Final proof. We now have all the tools necessary for the:

Proof of Theorem 4.16. We must verify each of the conditions from Definition 4.8.

Item (1) (bound on length of Ξ -chains) follows from Lemma 4.21(a), which implies that any chain $\text{Lk}(\Delta_1) \subsetneq \dots$ has length bounded by the number of possible types, which is 9.

Let us now discuss item (2) of the definition. The descriptions of the $\mathcal{C}(\Delta)$ from Lemmas 4.17, 4.18, 4.19, 4.20 yields that all $\mathcal{C}(\Delta)$ are hyperbolic, since each of them is either bounded or uniformly quasi-isometric to one of \hat{E} (which is hyperbolic by

Theorem 2.1), some Ξ^v (which is hyperbolic by Lemma 3.10), or \mathbb{R} . Moreover, any $\mathcal{C}(\Delta)$ is (uniformly) quasi-isometrically embedded in Y_Δ by Lemma 4.22.

Item (3) of the definition (common nesting) is precisely Lemma 4.21(b).

Finally, we show item (4) of the definition (fullness of links), which we recall for the convenience of the reader:

- If v, w are distinct non-adjacent vertices of $\text{Lk}(\Delta)$, for some simplex Δ of \mathcal{X} , contained in \mathcal{W} -adjacent maximal simplices, then they are contained in \mathcal{W} -adjacent simplices of the form $\Delta \star \Delta'$.

It suffices to consider simplices Δ of Ξ - and \mathcal{K} -type. Indeed, in all other cases (see Lemma 4.20), the vertices v, w under consideration are contained in the link of a simplex Δ' containing Δ where Δ' is of Ξ - or \mathcal{K} -type (as can be seen by enlarging Δ until its link is no longer a join; v and w are not \mathcal{X} -adjacent so they are contained in the same “side” of any join structure). Hence, once we deal with those cases, we know that there are suitable maximal simplices containing the larger simplex, whence Δ .

Consider first a simplex Δ of \mathcal{K} -type with vertices $s, v(s), w$. Consider distinct vertices t_1, t_2 (necessarily in \mathcal{K}) of $\text{Lk}(\Delta)$, and suppose that there are vertices $s_1, s_2 \in \mathcal{K}$ so that the maximal simplices $\sigma(s_1, t_1)$ and $\sigma(s_2, t_2)$ are connected in \mathcal{W} . There are two possibilities:

- $\bar{d}(M(s_1, t_1), M(s_2, t_2)) \leq 10R$. In this case, we have $\bar{d}(M(t_1), M(t_2)) \leq 10R$. In particular, in view of the second bullet in the definition of the edges of \mathcal{W} , we have that t_1, t_2 are contained, respectively, in the \mathcal{W} -connected maximal simplices $\Delta * t_1 = \sigma(s, t_1)$ and $\Delta * t_2 = \sigma(s, t_2)$.
- $s_1 = s_2$ and $\bar{d}(M(t_1), M(t_2)) \leq 10R$ (notice that $t_1 \neq t_2$ so that the “symmetric” case cannot occur). Again, we reach the same conclusion as above.

We can now consider a simplex Δ of Ξ -type with vertices $s, v(s)$. Consider vertices x_1, x_2 of $\text{Lk}(\Delta)$ that are not \mathcal{X} -adjacent but are contained in \mathcal{W} -adjacent maximal simplices. Furthermore, we can assume that x_1, x_2 are not in the link of a simplex of \mathcal{K} -type (the case we just dealt with) which contains Δ , since in that case we already know that there are suitable maximal simplices containing the larger simplex, whence Δ . Then, using the structure of $\text{Lk}(\Delta)$, we see that there are vertices $s_i, t_i \in \mathcal{K}$ so that:

- $x_i \in \{t_i, v(t_i)\}$,
- t_i and $v(t_i)$ all belong to $\text{Lk}(\Delta)$, and
- $\sigma(s_1, t_1)$ and $\sigma(s_2, t_2)$ are connected in \mathcal{W} .

In turn, the last bullet splits into two cases:

- $\bar{d}(M(s_1, t_1), M(s_2, t_2)) \leq 10R$. In this we have $\bar{d}(M(t_1), M(t_2)) \leq 10R$, so that x_1, x_2 are contained, respectively, in the \mathcal{W} -connected maximal simplices $\Delta * \{t_1, v(t_1)\} = \sigma(s, t_1)$ and $\Delta * \{t_2, v(t_2)\} = \sigma(s, t_2)$, so this case is fine.
- $s_1 = s_2$ and $\bar{d}(M(t_1), M(t_2)) \leq 10R$. But again we reach the same conclusion as before.

We now also have to check the existence of an action of Γ with the required properties. The action is constructed in Lemma 4.15, where all properties are checked except finiteness of the number of orbits of links of \mathcal{X} . The finitely many

possible types of links are listed in Lemmas 4.17 – 4.20, and for each type of simplex there are only finitely many orbits, so we are done. \square

5. QUASI-ISOMETRIC RIGIDITY

In this section, using the HHS structure, we prove a strong form of rigidity for the group Γ and the model space \bar{E} . Recall that \bar{E} is defined via a particular truncation \bar{D} of the Teichmüller disk D obtained by removing 1-separated horoballs. We say that such a truncation \bar{E} is an *allowable truncation* of E if Γ acts by isometries on it with cocompact quotient. Write $\text{Isom}(\Omega)$ and $\text{QI}(\Omega)$ for the isometry group and quasi-isometry group, respectively, of a metric space Ω . For \bar{E} , we write $\text{Isom}_{\text{fib}}(\bar{E}) \leq \text{Isom}(\bar{E})$ for the subgroup of isometries that map fibers to fibers.

Theorem 1.7. *There is an allowable truncation \bar{E} of E such that the natural homomorphisms $\text{Isom}_{\text{fib}}(\bar{E}) \rightarrow \text{Isom}(\bar{E}) \rightarrow \text{QI}(\bar{E}) \cong \text{QI}(\Gamma)$ are all isomorphisms, and $\Gamma \leq \text{Isom}(\bar{E}) \cong \text{QI}(\Gamma)$ has finite index.*

The proof is divided up into several steps which we outline here before getting into the details. The first step is to use the HHS structure to identify certain quasi-flats in \bar{E} , and prove that they are coarsely preserved by a quasi-isometry. The maximal quasi-flats are encoded by the strip bundles in \bar{E} , and using the preservation of quasi-flats, we show that a quasi-isometry further preserves strip bundles, and even sends all strip bundles for strips in any fixed direction to strip bundles in some other fixed direction. From there we deduce that a quasi-isometry sends fibers E_X within a bounded distance of some other fibers E_Y , and in fact induces a quasi-isometry between the fibers. Fixing attention on E_0 and further appealing to the structure of strip bundles, we show that a self quasi-isometry of \bar{E} induces a special type of quasi-isometry from E_0 to itself sending strips in a fixed direction within a uniformly bounded distance of strips in some other fixed direction. This quasi-isometry is promoted to a piecewise affine biLipschitz map from E_0 to itself, which we then show is in fact affine. This produces a homomorphism to the full affine group of E_0 , $\text{QI}(\bar{E}) \rightarrow \text{Aff}(E_0)$. Given an affine homeomorphism of E_0 , we construct an explicit fiber preserving isometry associated to it, which via the inclusions $\text{Isom}_{\text{fib}}(\bar{E}) \rightarrow \text{QI}(\bar{E})$ serves as a one-sided inverse. Finally, we prove that the homomorphism $\text{QI}(\bar{E}) \rightarrow \text{Aff}(E_0)$ is injective, hence the homomorphisms $\text{Isom}_{\text{fib}}(\bar{E}) \rightarrow \text{Isom}(\bar{E}) \rightarrow \text{QI}(\bar{E}) \rightarrow \text{Aff}(E_0)$ are all isomorphisms. The fact that Γ has finite index in $\text{Isom}_{\text{fib}}(\bar{E})$, and hence in $\text{Isom}(\bar{E})$, is straightforward using the cocompactness of the action of Γ and the singular structure.

5.1. HHS structure and quasi-flats. Denote by \mathfrak{S}_0 the set $\mathcal{V} \times \{0, 1\}$. We denote the element $(v, 0)$ by v^{qt} (for “quasi-tree”) and $(v, 1)$ by v^{ql} (for “quasi-line”).

We denote by \mathcal{F} the set of all strip bundles of \bar{E} , that is, subbundles with fiber a strip and base the horocycle corresponding to the direction of the strip. (Roughly, these are the flats of the peripheral graph manifolds.)

Proposition 5.1 (Properties of the HHS structure). *The HHS structure (\bar{E}, \mathfrak{S}) on \bar{E} coming from Theorem 4.16 has the following properties, for some $K \geq 1$.*

- (1) *The set of non- \sqsubseteq -maximal $Y \in \mathfrak{S}$ with $\text{diam}(\mathcal{C}(Y)) \geq 4$ is in bijection with \mathfrak{S}_0 . Under said bijection:*
- (2) *$\mathcal{C}(v^{qt})$ is (K, K) -quasi-isometric to a quasi-tree with at least two points at infinity, and $\mathcal{C}(v^{ql})$ is (K, K) -quasi-isometric to a line;*

- (3) For all $v \in \mathcal{V}$, we have $v^{qt} \perp v^{ql}$;
- (4) For all adjacent $v, w \in \mathcal{V}$, we have $w^{ql} \perp v^{ql}$ and $w^{ql} \sqsubseteq v^{qt}$;
- (5) All pairs of elements of \mathfrak{S}_0 that do not fall into the aforementioned cases are transverse;
- (6) For each adjacent $v, w \in \mathcal{V}$ there is $F \in \mathcal{F}$ so $\pi_{v^{qt}}(F)$ and $\pi_{w^{ql}}(F)$ are K -coarsely dense, and $\pi_Y(F)$ has diameter at most K for all $Y \neq v^{ql}, w^{ql}$.

Proof. The second paragraph of the proof of Theorem 4.16 implies $\mathcal{C}(v^{qt})$ is quasi-isometric to the quasi-tree Ξ^v (with at least two points at infinity by Lemma 3.10) and $\mathcal{C}(v^{ql})$ is quasi-isometric to the quasi-line \mathcal{K}^v , and that these are the only non-maximal elements of diameter at least 4. This proves (1) and (2). In view of the combinatorial description of orthogonality and nesting from Definition 4.9, properties (3)-(5) boil down to combinatorial properties of \mathcal{X} that are straightforward to check. For example, regarding property (3) note that two (equivalence classes of) simplices are orthogonal if their links form a join. The links of the simplices corresponding to v^{qt} and v^{ql} are $\text{Lk}_{\Xi}(v)$ and \mathcal{K}^v (see Lemmas 4.18 and 4.19), which indeed form a join.

Regarding property (6), first of all the projections in the HHS structure on \bar{E} are obtained composing the quasi-isometry $\bar{E} \rightarrow \mathcal{W}$ from Lemma 4.15 and the projections defined in Definition 4.10 (roughly, those are closest-point projections in the complement of saturations).

The required strip bundle is the intersection $\Theta^v \cap \Theta^w$, which under the quasi-isometry of Lemma 4.15 corresponds to the set of all maximal simplices of \mathcal{W} of the form $\sigma(s, t)$ for $s \in \mathcal{K}^v$, $t \in \mathcal{K}^w$. In view of the description of the π_Y from Definition 4.10, the coarse density claim follows since the union of the simplices described above contains the links of the simplices corresponding to v^{ql} and w^{ql} , which are \mathcal{K}^v and \mathcal{K}^w .

Regarding the boundedness claim, it can be checked case-by-case that the set of simplices described above gives a bounded set of Y_{Δ} for $[\Delta] \neq v^{ql}, w^{ql}$ (for example, note that said set is bounded if the saturation of Δ does not intersect $\mathcal{K}^v \cup \mathcal{K}^w$, or if it does not contain v or w). This implies boundedness of the projections since the projections are coarsely Lipschitz; this follows from Theorem 4.16 since the projection maps of an HHS are required to be coarsely Lipschitz. \square

From now on we identify \mathfrak{S}_0 with the set of all $Y \in \mathfrak{S}$ with $\text{diam}(\mathcal{C}(Y)) \geq 4$ as in Proposition 5.1. Notice that the maximal number of pairwise orthogonal elements of \mathfrak{S}_0 is 2. Therefore, a *complete support set* as in [BHS21, Definition 5.1] is just a pair of orthogonal elements of \mathfrak{S}_0 .

Let \mathcal{H} be the set of pairs (Y, p) where $Y \in \mathfrak{S}$ and $p \in \partial\mathcal{C}(Y)$ with $Y = v^{ql}$ for some $v \in \mathcal{V}$. We say that two such pairs (Y, p) and (W, q) are orthogonal if Y and W are. Any element $\sigma = (Y, p) \in \mathcal{H}$ comes with a quasi-geodesic ray h_{σ} in \bar{E} , as in [BHS21, Definition 5.3], so that $\pi_Y \circ h_{\sigma}$ is a quasi-geodesic in $\mathcal{C}(Y)$ and $\pi_W(h_{\sigma})$ is bounded for all $W \neq Y$.

We recall that given subsets A and B of a metric space X , we say that the subset C of X is the coarse intersection of A and B if for every sufficiently large R we have that $N_R(A) \cap N_R(B)$ lies within finite Hausdorff distance of C . If the coarse intersection of two subsets exists, then it is well-defined up to finite Hausdorff distance.

Lemma 5.2. *Let $\phi: \bar{E} \rightarrow \bar{E}$ be a quasi-isometry. Then there is a bijection $\phi_{\mathcal{H}}: \mathcal{H} \rightarrow \mathcal{H}$ preserving orthogonality and so that $d_{\text{Haus}}(\phi(h_{\sigma}), h_{\phi_{\mathcal{H}}(\sigma)}) < \infty$.*

Proof. Let \mathcal{H}' be the set of pairs (Y, p) where $Y \in \mathfrak{S}$ and $p \in \partial C(Y)$, without the restriction on Y .

The lemma with \mathcal{H}' replacing \mathcal{H} would follow directly from [BHS21, Theorem 5.7], except that the HHS structure on \bar{E} does not satisfy one of the 3 required assumptions, namely Assumption 2 (while it does satisfy Assumption 1 by parts (1) and (2) of Proposition 5.1, and it also satisfies Assumption 3 since there are no 3 pairwise orthogonal elements of \mathfrak{S}_0 , by parts (3)-(5) of Proposition 5.1).

Inspecting the proof of [BHS21, Theorem 5.7], we see that Assumption 2 is used in two places.

The first one is to define the map $\phi_{\mathcal{H}}$ on a certain pair $\sigma = (Y, p) \in \mathcal{H}'$. The argument applies verbatim if Y satisfies Assumption 2, that is, if and only if Y is the intersection of 2 complete support sets. This is the case if $Y = v^{qt}$ for some $v \in \mathcal{H}$, that is, if $\sigma \in \mathcal{H}$. Therefore, one can use that argument to define a map $\phi_{\mathcal{H}}: \mathcal{H} \rightarrow \mathcal{H}'$. What is more, the image of $\phi_{\mathcal{H}}$ needs to be contained in \mathcal{H} . This can be seen from the fact that $h_{\phi_{\mathcal{H}}(\sigma)}$ for $\sigma \in \mathcal{H}$ arises as a coarse intersection of *standard orthants*, which are, essentially, products of rays h_{σ} , see [BHS21, Definition 4.1] for the precise definition. Notice that [BHS21, Lemma 4.11] says, roughly, that coarse intersections of standard orthants are the expected sub-products. Hence, the failure of Assumption 2 for $Y = v^{qt}$ implies that $h_{(Y,p)}$ cannot be a coarse intersection of standard orthants, and therefore $\phi_{\mathcal{H}}(\sigma)$ for $\sigma \in \mathcal{H}$ also needs to lie in \mathcal{H} .

The second place where Assumption 2 is mentioned in [BHS21, Theorem 5.7] is the proof that $\phi_{\mathcal{H}}$ preserves orthogonality. There the assumption is used to say that certain quasi-geodesic rays are of the form h_{σ} . Such quasi-geodesic rays arise as coarse intersections of standard orthants, so, as mentioned above, they need to be of the form h_{σ} for $\sigma \in \mathcal{H}$, hence Assumption 2 is not actually needed there.

Thus, the arguments in the proof of [BHS21, Theorem 5.7] give the lemma. \square

Lemma 5.3. *For every K there exists C so that the following holds. Let $\phi: \bar{E} \rightarrow \bar{E}$ be a (K, K) -quasi-isometry. Then there is a bijection $\phi_{\mathcal{F}}: \mathcal{F} \rightarrow \mathcal{F}$ so that $d_{\text{Haus}}(\phi(F), \phi_{\mathcal{F}}(F)) \leq C$ for all $F \in \mathcal{F}$.*

Proof. Let p^{\pm} be the two points at infinity of $\mathcal{C}(v^{qt})$ for some $v \in \mathcal{V}$. We claim that there exists $w \in \mathcal{V}$ so that, for q^{\pm} the points at infinity of $\mathcal{C}(w^{qt})$, we have $\phi_{\mathcal{H}}((v^{qt}, p^{\pm})) = (w^{qt}, q^{\pm})$, up to relabeling. We use that $\phi_{\mathcal{H}}$ preserves orthogonality to show this. Let $u_1, u_2 \in \mathcal{V}$ be distinct and adjacent to v , and let r_1^{\pm}, r_2^{\pm} be the points at infinity of $\mathcal{C}(u_1^{qt}), \mathcal{C}(u_2^{qt})$. Then (v^{qt}, p^{\pm}) are the only elements of \mathcal{H} that are orthogonal to all the (u_i^{qt}, r_i^{\pm}) . Since $\phi_{\mathcal{H}}$ preserves orthogonality, we see that $\phi_{\mathcal{H}}((v^{qt}, p^{\pm}))$ are both orthogonal to the same 4 distinct elements of \mathcal{H} with the property that no pair of them is orthogonal. This is easily seen to imply that $\phi_{\mathcal{H}}((v^{qt}, p^{\pm}))$ must be of the form (w, q^{\pm}) , since said 4 elements need to be associated to at least 2 distinct vertices of \mathcal{V} . This shows the claim.

In view of the claim, we see that [BHS21, Lemma 5.9] applies. (We note that Assumption 2 in said Lemma is only needed to have the map from [BHS21, Theorem 5.7], but our map from Lemma 5.2 has the same defining properties, just with a smaller domain and range.)

The *standard flats* in [BHS21, Lemma 5.9] coarsely coincide with the elements of \mathcal{F} in view of Proposition 5.1(6) (compare with [BHS21, Definition 4.1]) so the lemma follows. \square

Denote by \mathcal{S} the collection of all strips in E_0 , and for $A \in \mathcal{S}$ denote by $\alpha(A) \in \mathcal{P}$ the direction of A . Similarly, for $F \in \mathcal{F}$ we denote $\alpha(F) \in \mathcal{P}$ the direction of the strip defining F .

Proposition 5.4. *Given $K \geq 1$, there exists $C \geq 0$ so that if $\phi: \bar{E} \rightarrow \bar{E}$ is a (K, K) -quasi-isometry, then for all $X \in \bar{D}$, there exists $Y \in \bar{D}$ so that the Hausdorff distance between $\phi(E_X)$ and E_Y is at most C . In particular, $d_{\text{Haus}}(E_0, \phi(E_0)) < \infty$. Moreover, there are bijections $\phi_{\mathcal{P}}: \mathcal{P} \rightarrow \mathcal{P}$ and $\phi_{\mathcal{S}}: \mathcal{S} \rightarrow \mathcal{S}$ so that:*

- (1) $d_{\text{Haus}}(\phi(\partial\mathcal{B}_{\alpha}), \partial\mathcal{B}_{\phi_{\mathcal{P}}(\alpha)}) \leq C$ for each $\alpha \in \mathcal{P}$.
- (2) $\alpha(\phi_{\mathcal{S}}(A)) = \phi_{\mathcal{P}}(\alpha(A))$ and $d_{\text{Haus}}(\phi(A), \phi_{\mathcal{S}}(A)) < \infty$ for all $A \in \mathcal{S}$,

Proof. First, note that fibers are quantitatively coarse intersections of the sets $\partial\mathcal{B}_{\alpha}$, in the sense that exists a function $f: \mathbb{R} \rightarrow \mathbb{R}$ and $t_0 \geq 0$ such that

- for any $X \in \bar{D}$ and any $t \geq t_0$ there are two distinct $\partial\mathcal{B}_{\alpha}$ whose t -neighborhoods intersect in a set within Hausdorff distance $f(t)$ of E_X ;
- for $t \geq t_0$, if the t -neighborhoods of two distinct $\partial\mathcal{B}_{\alpha}$ intersect, then this intersection lies within Hausdorff distance $f(t)$ of a fiber.

This follows via the bundle-map $\pi: \bar{E} \rightarrow \bar{D}$ and the corresponding relationship between neighborhoods of distinct horocycles $\partial\mathcal{B}_{\alpha}$ in \bar{D} .

We next make three preliminary observations. Firstly, for each $\alpha \in \mathcal{P}$ the set $\partial\mathcal{B}_{\alpha}$ is the union of all $F \in \mathcal{F}$ with $\alpha(F) = \alpha$. Secondly, if $F_1, F_2 \in \mathcal{F}$ have $\alpha(F_1) \neq \alpha(F_2)$ then the coarse intersection of F_1 and F_2 is bounded. Indeed, F_i is contained in $\partial\mathcal{B}_{\alpha(F_i)}$, and the coarse intersection of these $\partial\mathcal{B}_{\alpha(F_i)}$ is some (or really, any) fiber E_X . Since the coarse intersection of F_i with E_X is a strip in the corresponding direction, and strips in different directions have bounded coarse intersection, the claim follows. Thirdly, observe that $F_1, F_2 \in \mathcal{F}$ have $\alpha(F_1) = \alpha(F_2)$ if and only if there is a chain of elements in \mathcal{F} from F_1 to F_2 so that consecutive elements have unbounded coarse intersection. The “if” part follows from the previous observation, while the “only if” follows from the fact that elements of \mathcal{F} corresponding to adjacent edges of some T_{α} have unbounded coarse intersection.

In view of all this and Lemma 5.3, we see that for each α there exists a (necessarily unique) $\phi_{\mathcal{P}}(\alpha) \in \mathcal{P}$ so that $\phi(\partial\mathcal{B}_{\alpha})$ and $\partial\mathcal{B}_{\phi_{\mathcal{P}}(\alpha)}$ have finite Hausdorff distance. In fact, the distance is uniformly bounded by the constant C , depending only on K , coming from Lemma 5.3. This is how we define $\phi_{\mathcal{P}}$.

Now for any $X \in \bar{D}$, we may choose $\alpha_1, \alpha_2 \in \mathcal{P}$ so that the fiber E_X has Hausdorff distance at most $f(t_0)$ from the intersection of the t_0 -neighborhoods of $\partial\mathcal{B}_{\alpha_i}$. Thus there is some uniform $t'_0 \geq t_0$, again depending only on K , so that $\phi(E_X)$ has Hausdorff distance at most t'_0 from the intersection of the t'_0 -neighborhoods of $\partial\mathcal{B}_{\phi_{\mathcal{P}}(\alpha_i)}$; further, as mentioned above, this intersection of t'_0 -neighborhoods has Hausdorff distance at most $f(t'_0)$ to some fiber E_Y , as claimed.

Finally, we define $\phi_{\mathcal{S}}$ via the bijection $\mathcal{F} \leftrightarrow \mathcal{S}$ between strip bundles and strips in E_0 . That is, if $A \in \mathcal{S}$ corresponds to $F \in \mathcal{F}$, then $\phi_{\mathcal{S}}(A)$ is the strip corresponding to $\phi_{\mathcal{F}}(F)$. Since A is the coarse intersection of F with E_0 , the desired properties for $\phi_{\mathcal{S}}$ then follow from the facts that $\alpha_{\phi_{\mathcal{F}}(F)} = \phi_{\mathcal{P}}(\alpha(F))$ and that $\phi_{\mathcal{F}}(F)$ lies within finite Hausdorff distance of $\phi(F)$. \square

5.2. From $\text{QI}(\bar{E})$ to $\text{QI}(E_0)$. The next step is to construct a homomorphism $\text{QI}(\bar{E}) \rightarrow \text{QI}(E_0)$ by associating a quasi-isometry of E_0 to each quasi-isometry of \bar{E} (see Lemma 5.7). This step requires some preliminaries which we now explain.

To distinguish between two relevant notions of properness, we will call a map $f: X \rightarrow Y$ between metric spaces *topologically proper* if it is continuous and preimages of compact sets are compact, and *metrically proper* if there exist diverging functions $\rho_-, \rho_+ : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ (which we will call properness functions) such that for all $x, y \in X$ we have

$$\rho_-(d_X(x, y)) \leq d_Y(f(x), f(y)) \leq \rho_+(d_X(x, y)).$$

(Both types of maps are just referred to as “proper” in the appropriate contexts, but neither notion implies the other.)

For $R > 0$ and $X \in \bar{D}$, we endow $N_R(E_X)$ with the restriction of the metric of \bar{E} , while E_X is endowed with its path metric. Then the restriction of $f_X: \bar{E} \rightarrow E_X$ to $N_R(E_X)$ is metrically proper. Indeed, f_X is topologically proper and equivariant with respect to a group acting cocompactly. Note that the properness functions here can be taken independently of the fiber X (once we fix R) because there is also a cocompact action on \bar{D} . We also note the following lemma.

Lemma 5.5. *A metrically proper coarsely surjective map between geodesic metric spaces is a quasi-isometry. Moreover, the quasi-isometry constants depend only on the properness functions and the coarse surjectivity constant.*

This follows from standard arguments. First, a metrically proper map from a geodesic metric space is coarsely Lipschitz (the proof involves subdividing geodesics into segments of length at most 1, each of which has bounded image). Also, coarse surjectivity allows one to construct a quasi-inverse of the map, which is furthermore metrically proper. As above, the quasi-inverse is coarsely Lipschitz, and we conclude since a coarsely Lipschitz map with a coarsely Lipschitz quasi-inverse is a quasi-isometry.

Given any quasi-isometry $\phi: \bar{E} \rightarrow \bar{E}$ and $X \in \bar{D}$, define $\nu_\phi^X: E_X \rightarrow E_X$ to be $\nu_\phi^X = f_X \circ \phi|_{E_X}$. In the case of the base fiber $X = X_0$ we denote this $\nu_\phi = \nu_\phi^{X_0}$. When ϕ is understood, we also write $\nu^X = \nu_\phi^X$ and $\nu = \nu_\phi$.

Lemma 5.6. *For any (K, K) -quasi-isometry $\phi: \bar{E} \rightarrow \bar{E}$ and $X \in \bar{D}$, the map $\nu_\phi^X: E_X \rightarrow E_X$ is a (K', K') -quasi-isometry, where K' depends only on K and $d_{\text{Haus}}(E_X, \phi(E_X))$. Furthermore, for any $A \in \mathcal{S}$, $d_{\text{Haus}}(\nu_\phi^X(A), \phi_S(A)) < \infty$.*

Proof. First note that the restriction of ϕ to E_X is metrically proper, since the path metric on E_X and the restricted metric from \bar{E} are coarsely equivalent (that is, the identity on E_X is a metrically proper map between these metric spaces). Next let $R = d_{\text{Haus}}(E_X, \phi(E_X))$, which is finite by Proposition 5.4, and note that the restriction $f_X|_{N_R(E_X)}: N_R(E_X) \rightarrow E_X$ is also metrically proper. Therefore the composition $\nu_\phi^X = (f_X|_{N_R(E_X)}) \circ (\phi|_{E_X})$ is metrically proper and, moreover, the properness functions depend only on K, R and not on the fiber E_X .

By [KL12, Theorem 3.8] and the fact that E_X is uniformly quasi-isometric to \mathbb{H}^2 , any metrically proper map of E_X to itself is coarsely surjective and, moreover, the coarse surjectivity constant depends only on the properness functions. Therefore ν_ϕ^X is coarsely surjective and a uniform quasi-isometry by Lemma 5.5.

Regarding the claim about A , this follows from Proposition 5.4(2) and the fact that f_X moves each point of $\phi(A) \subseteq N_R(E_X)$ at most R away. \square

Lemma 5.7. *The assignment $\phi \mapsto \nu_\phi$, for any quasi-isometry $\phi: \bar{E} \rightarrow \bar{E}$, well-defines a homomorphism $\mathcal{A}_0: \text{QI}(\bar{E}) \rightarrow \text{QI}(E_0)$.*

Proof. Given any quasi-isometry $\phi: \bar{E} \rightarrow \bar{E}$ and $x \in E_0$ we have

$$d(\phi(x), \nu_\phi(x)) = d(\phi(x), f_0(\phi(x))) \leq d_{\text{Haus}}(\phi(E_0), E_0)$$

The right hand side is finite by Proposition 5.4, so the left hand side is bounded, independent of x . From this, the triangle inequality, and the uniform metric properness of the inclusion of E_0 into \bar{E} , it easily follows that if ϕ and ϕ' are bounded distance, then so are ν_ϕ and $\nu_{\phi'}$. Therefore the assignment $\phi \mapsto \nu_\phi$ descends to a well-defined function $\mathcal{A}_0: \text{QI}(\bar{E}) \rightarrow \text{QI}(E_0)$.

To see that \mathcal{A}_0 is a homomorphism, suppose ϕ, ϕ' are (K, K) -quasi-isometries of \bar{E} . Then from the inequality above, for all $x \in E_0$ we have

$$d(\phi' \circ \phi(x), \phi' \circ \nu_\phi(x)) \leq Kd(\phi(x), \nu_\phi(x)) + K \leq Kd_{\text{Haus}}(\phi(E_0), E_0) + K.$$

The left-hand side is thus uniformly bounded, independent of x . From this, the triangle inequality, and Proposition 5.4, it follows that $d_{\text{Haus}}(\phi' \circ \phi(E_0), E_0)$ and $d_{\text{Haus}}(\phi' \circ \nu_\phi(E_0), E_0)$ are bounded by some constant $r > 0$. Then for all $x \in E_0$,

$$d(\nu_{\phi' \circ \phi}(x), \nu_{\phi'} \circ \nu_\phi(x)) = d(f_0(\phi' \circ \phi(x)), f_0(\phi' \circ \nu_\phi(x))) \leq e^r d(\phi' \circ \phi(x), \phi' \circ \nu_\phi(x)).$$

Combining this with the previous inequality, we see that the quantity on the right, and hence the left, is uniformly bounded above, independent of x . Therefore $\nu_{\phi' \circ \phi}$ and $\nu_{\phi'} \circ \nu_\phi$ are bounded distance apart and \mathcal{A}_0 is a homomorphism. \square

5.3. From quasi-isometries to affine homeomorphisms. The flat metric q on E_0 determines an associated affine group $\text{Aff}(E_0)$, and we observe that if $\phi \in \Gamma$ is an element of the extension group (which is an isometry of \bar{E} , and so also a quasi-isometry), we have $\nu_\phi \in \text{Aff}(E_0)$. The next step in the proof of rigidity is the following.

Proposition 5.8. *For any quasi-isometry $\phi: \bar{E} \rightarrow \bar{E}$, the quasi-isometry ν_ϕ is uniformly close to a unique element $\nu_\phi^a \in \text{Aff}(E_0)$.*

The proof of the proposition will take place over the remainder of this subsection. Before getting to the proof, however, we note a useful corollary. Two quasi-isometries ϕ_1, ϕ_2 that are a bounded distance apart have ν_{ϕ_1} and ν_{ϕ_2} a bounded distance apart, and so by the uniqueness $\nu_{\phi_1}^a = \nu_{\phi_2}^a$. Thus we have the following.

Corollary 5.9. *The map $[\phi] \mapsto \nu_\phi^a$ defines a homomorphism $\mathcal{A}: \text{QI}(\bar{E}) \rightarrow \text{Aff}(E_0)$. Moreover, the homomorphism $\mathcal{A}_0: \text{QI}(\bar{E}) \rightarrow \text{QI}(E_0)$ from Lemma 5.7 factors as the composition of \mathcal{A} with the natural inclusion $\text{Aff}(E_0) \rightarrow \text{QI}(E_0)$.*

Fix a triangulation \mathfrak{t} of X_0 so that the vertex set is the set of cone points and all triangles are Euclidean triangles (that is, they are images of triangles by maps that are locally isometric and injective on the interior; see e.g. [DDLS21, Lemma 3.4]). Moreover, we assume that all saddle connections in some direction α_0 appear as edges of the triangulation. Lift \mathfrak{t} to a triangulation $\tilde{\mathfrak{t}}$ of E_0 . By assumption, all saddle connections in E_0 in direction α_0 are edges of $\tilde{\mathfrak{t}}$, and the complement of the union of this subset is a union of all (interiors of) strips in direction α_0 .

Lemma 5.10. *Given a quasi-isometry ϕ , there is a biLipschitz homeomorphism $\nu_\phi^a: E_0 \rightarrow E_0$ a bounded distance from ν_ϕ so that ν_ϕ^a restricts to an affine map on*

each triangle of $\tilde{\mathfrak{k}}$. Moreover, if an edge δ of $\tilde{\mathfrak{k}}$ has direction $\alpha \in \mathcal{P}$, then $\nu_\phi^\alpha(\delta)$ has direction $\phi_{\mathcal{P}}(\alpha)$.

We will later prove that ν_ϕ^α is in fact *globally* affine, justifying the notation.

Proof. Given $\nu = \nu_\phi: E_0 \rightarrow E_0$, let $\partial\nu: S_\infty^1 \rightarrow S_\infty^1$ be the restriction of the extension to the Gromov boundary S_∞^1 of E_0 . The space \mathcal{G} of (unordered) pairs of distinct points in S_∞^1 is precisely the space of endpoint-pairs at infinity of unoriented biinfinite geodesics (up to the equivalence relation of having finite Hausdorff distance). The map $\partial\nu$ induces a map $\partial\nu_*: \mathcal{G} \rightarrow \mathcal{G}$.

Let $\mathcal{G}^* \subset \mathcal{G}$ be the closure of the set of endpoint-pairs at infinity of non-singular geodesics (i.e. geodesics that miss every cone point). Observe that all geodesics in a given strip have the same pair of endpoints, and any geodesic with that pair of endpoints is contained in the strip. Given a strip, we are therefore justified in referring to the *pair of endpoints of the strip*.

It follows from the description of geodesics with endpoints in \mathcal{G}^* (see [BL18, Proposition 2.4]) together with the *Veech Dichotomy* (see e.g. [MT02]), that for any $\{\xi, \zeta\} \in \mathcal{G}^*$, either $\{\xi, \zeta\}$ are the endpoints of a strip, or endpoints of a geodesic meeting at most one cone point.

According to Proposition 5.4, for any strip $A \in \mathcal{S}$, the strip $\phi_{\mathcal{S}}(A)$ has finite Hausdorff distance to $\phi(A)$, and hence it also has finite Hausdorff distance to $\nu(A)$. Since $\phi_{\mathcal{S}}$ is a bijection, this means that the homeomorphism $\partial\nu_*$ sends the dense subset of \mathcal{G}^* consisting of endpoint of strips onto itself, hence $\nu_*(\mathcal{G}^*) = \mathcal{G}^*$. From this and [BL18, Proposition 4.1] (see also [DELS18, Proposition 11]), it follows that there is a bijection

$$\phi_{\Sigma_0}: \Sigma_0 \rightarrow \Sigma_0$$

from the set of cone points Σ_0 of E_0 to itself with the following property. If $\gamma \subset E_0$ is a geodesic or strip containing $x \in \Sigma_0$ with endpoints $\{\xi, \zeta\} \in \mathcal{G}^*$, then $\partial\nu_*(\{\xi, \zeta\})$ are the endpoints of a geodesic containing $\phi_{\Sigma_0}(x)$. Given $x \in \Sigma_0$ consider any two geodesics γ_1 and γ_2 with endpoints in \mathcal{G}^* (not necessarily contained in strips) passing through x making an angle at least $\pi/2$ with each other. We note that $\nu(x)$ is contained in $\nu(\gamma_1)$ and $\nu(\gamma_2)$, and is thus some uniform distance $r > 0$ to both of their geodesic representatives. Since γ_1 and γ_2 meet at angle at least $\pi/2$, the r -neighborhoods of the geodesic representatives of $\nu(\gamma_1)$ and $\nu(\gamma_2)$ intersect in a uniformly bounded diameter set, which contains $\phi_{\Sigma_0}(x)$. Therefore, $\phi_{\Sigma_0}(x)$ is uniformly close to $\nu(x)$, for all $x \in \Sigma_0$.

From the properties of ϕ_{Σ_0} described above, we see that if $x \in \Sigma_0$ is contained in a strip A , then $\phi_{\Sigma_0}(x)$ is contained in the strip $\phi_{\mathcal{S}}(A)$. For any saddle connection δ in some direction $\alpha \in \mathcal{P}$ between a pair of points $x, y \in \Sigma_0$, there is a unique pair of strips A_1, A_2 , also in direction α , that contain δ . Since $\phi_{\Sigma_0}(x), \phi_{\Sigma_0}(y)$ are contained in $\phi_{\mathcal{S}}(A_1)$ and $\phi_{\mathcal{S}}(A_2)$, it follows that there is a unique saddle connection with endpoints $\phi_{\Sigma_0}(x), \phi_{\Sigma_0}(y)$. For any strip A the saddle connections whose union makes up one of its boundary components is determined by a collection of strips meeting A in the given saddle connections. Considering the cyclic ordering of the endpoints of these strips (and those of A) on S_∞^1 , and the fact that $\partial\nu$ is a homeomorphism, it follows that ϕ_{Σ_0} maps the *ordered* set of cone points along each boundary component of the strip A by an order preserving (or reversing) bijection to the ordered set of cone points along the boundary components of $\phi_{\mathcal{S}}(A)$.

We can now extend the map ϕ_{Σ_0} to a map $\nu_\phi^a: E_0 \rightarrow E_0$ using $\tilde{\mathfrak{t}}$ as follows. First, recall that any edge of $\tilde{\mathfrak{t}}$ is a saddle connection δ connecting two points $x, y \in \Sigma_0$. By the previous paragraph, there is a saddle connection δ' connecting $\phi_{\Sigma_0}(x)$ and $\phi_{\Sigma_0}(y)$, and we define ν_ϕ^a on δ so that it maps δ by an affine map to δ' extending ϕ_{Σ_0} on the endpoints. This defines ν_ϕ^a on the 1-skeleton, $\tilde{\mathfrak{t}}^1$, and since ϕ_{Σ_0} is a bounded distance from $\nu|_{\Sigma_0}$, it follows that $\nu_\phi^a|_{\tilde{\mathfrak{t}}^1}$ is a bounded distance from $\nu|_{\tilde{\mathfrak{t}}^1}$.

By our assumptions on $\tilde{\mathfrak{t}}$, there is a subset of the edges of $\tilde{\mathfrak{t}}$ whose union is precisely the union of boundaries of all strips in direction α_0 . The order preserving (or reversing) property described above for the cone points along the boundary of a strip, together with Proposition 5.4, implies that for any boundary component of any strip A in direction α_0 , ν_ϕ^a restricted to its boundary components is a homeomorphism onto the boundary components of $\phi_S(A)$. Furthermore, since the sides of any triangle of $\tilde{\mathfrak{t}}$ are contained in such a strip A , the ν_ϕ^a -image of the sides are contained in $\phi_S(A)$. We can now extend ν_ϕ^a over the triangles by the unique affine map extending the map on their sides.

Since disjoint strips map to disjoint strips, the map ν_ϕ^a is a homeomorphism. By construction, any edge in direction α is sent to an edge in direction $\phi_{\mathcal{P}}(\alpha)$. Since $\tilde{\mathfrak{t}}$ projects to \mathfrak{t} , there are only finitely many directions that the sides of a triangle can lie in and so finitely many isometry types of triangles. Each of these finitely many isometry types maps by an affine map to only finitely many types of triangles in the image (because the direction of the images of sides are determined by $\phi_{\mathcal{P}}$), and therefore these affine maps are uniformly biLipschitz. Therefore, ν_ϕ^a is biLipschitz, completing the proof. \square

To show that ν_ϕ^a is affine, we analyze the effect of using it to conjugate the action of $\pi_1 S$ on E_0 .

Lemma 5.11. *The action of $\pi_1 S$ on E_0 obtained by conjugating the isometric action by ν_ϕ^a is again an isometric action.*

Before proving the lemma, we use it to prove the proposition.

Proof of Proposition 5.8. By Lemma 5.11, $\Lambda = \nu_\phi^a \pi_1 S (\nu_\phi^a)^{-1}$ acts by isometries, and ν_ϕ^a descends to a homeomorphism $\mu_\phi^a: S \rightarrow E_0/\Lambda$ and is biLipschitz with respect to descent to S and E_0/Λ of q . Since ν_ϕ^a and ν are a bounded distance, they have the same boundary maps. Since $\partial(\nu_\phi^a)_* = \partial\nu_*$ maps \mathcal{G}^* to \mathcal{G}^* , the *Current Support Theorem* of [DELS18] (and its proof) implies that the descent of $\mu_\phi^a: (S, q) \rightarrow (E_0/\Lambda, q)$ is affine. Therefore ν_ϕ^a is an affine map which is a bounded distance from $\nu = \nu_\phi$, as required.

Uniqueness follows from the fact that no two distinct affine maps are a bounded distance apart. \square

Proof of Lemma 5.11. We need to show that for all $g \in \pi_1 S$, the map

$$\nu_\phi^a \circ g \circ (\nu_\phi^a)^{-1}: E_0 \rightarrow E_0$$

is an isometry. For this, fix a triangle τ of $\tilde{\mathfrak{t}}$ and consider the restriction to $\nu_\phi^a(\tau)$. Let $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{P}$ be the directions of the sides. Setting $\alpha'_i = \phi_{\mathcal{P}}(\alpha_i)$, for $i = 1, 2, 3$, Lemma 5.10 implies that the directions of the sides of $\nu_\phi^a(\tau)$ are $\alpha'_1, \alpha'_2, \alpha'_3$. The action of $\pi_1 S$ on E_0 is by isometries, but it also preserves parallelism (i.e. each element induces the identity on $\mathbb{P}^1(q)$). Therefore, for any $g \in \pi_1 S$, the directions

of the sides of $g(\tau)$ are also $\alpha_1, \alpha_2, \alpha_3$, and by Lemma 5.10 again, it follows that the sides of $\nu_\phi^a(g(\tau))$ are $\alpha'_1, \alpha'_2, \alpha'_3$.

For any $g \in \pi_1 S$, since ν_ϕ^a is affine on τ , the composition $\nu_\phi^a \circ g \circ (\nu_\phi^a)^{-1}$ is also affine on $\nu_\phi^a(\tau)$. On the other hand, it also preserves the directions of the sides, $\alpha'_1, \alpha'_2, \alpha'_3$. Therefore, the restriction of $\nu_\phi^a \circ g \circ (\nu_\phi^a)^{-1}$ is a Euclidean similarity. Triangles of $\nu_\phi^a(\tilde{t})$ that share a side are scaled by the same factor by the similarity $\nu_\phi^a \circ g \circ (\nu_\phi^a)^{-1}$ in each triangle (since this is the scaling factor on the shared side). Therefore, the similarities agree along edges, and hence $\nu_\phi^a \circ g \circ (\nu_\phi^a)^{-1}$ defines a global similarity of E_0 .

So, the action of $\pi_1 S$ on E_0 obtained by conjugating by ν_ϕ^a is an action by similarities. To see that the action is by isometries, suppose that for some element $g \in \pi_1 S$, the similarity $g_0 = \nu_\phi^a \circ g \circ (\nu_\phi^a)^{-1}$ scales the metric some number $\lambda \neq 1$. Taking the inverse if necessary, we can assume $\lambda < 1$. Fix any $x \in E_0$ and observe that $d_q(g_0(x), g_0^2(x)) = \lambda d_q(x, g_0(x))$, where d_q is the distance function on E_0 determined by q . Iterating this, it follows that

$$d_q(x, g_0^n(x)) \leq \sum_{k=1}^n d_q(g_0^{k-1}(x), g_0^k(x)) = \sum_{k=1}^n \lambda^{k-1} d_q(x, g_0(x)) \leq d_q(x, g_0(x)) \sum_{k=1}^{\infty} \lambda^{k-1}.$$

Since the right-hand side is a convergent geometric series, it follows that $\{g_0^n(x)\}_{n=1}^{\infty}$ is a Cauchy sequence. On the other hand, this sequence exits every compact set (since g_0 is an infinite order element of $\pi_1 S$), and since q is a complete metric on E_0 , thus we obtain a contradiction. Therefore, the conjugation action of $\pi_1 S$ is by isometries. \square

5.4. Injectivity of \mathcal{A} . Our next goal is to prove that $\mathcal{A}: \text{QI}(\bar{E}) \rightarrow \text{Aff}(E_0)$, the homomorphism from Corollary 5.9, is injective.

In preparation, it will be useful to have the following general fact about quasi-isometries of hyperbolic spaces, whose proof we sketch for convenience of the reader:

Lemma 5.12. *For each K, C, δ there exists R so that the following holds. Suppose that Z is δ -hyperbolic and that each $z \in Z$ lies within δ of all three sides of a nondegenerate ideal geodesic triangle. Let $f: Z \rightarrow Z$ be a (K, C) -quasi-isometry that lies within finite distance of the identity. Then f lies within distance R of the identity.*

Proof. Since f is within bounded distance of the identity, its extension $\partial f: \partial Z \rightarrow \partial Z$ is the identity. Hence if $z \in Z$ and Δ is an ideal geodesic triangle as in the statement, then $f(\Delta)$ is a (K, C) -quasigeodesic ideal triangle with the same endpoints as Δ . By the Morse lemma, there is a constant $\kappa = \kappa(K, C, \delta) > 0$ such that $f(z)$ lies within κ of the three quasi-geodesic sides of $f(\Delta)$, and these sides in turn lie within κ of the sides of Δ . Thus z and $f(z)$ both lie within $2\kappa + \delta$ of all three sides of Δ . Since the set of points within $2\kappa + \delta$ of all three sides of a nondegenerate geodesic triangle in a δ -hyperbolic space has diameter bounded in terms of κ and δ , we see that $d_Z(z, f(z))$ is bounded solely in terms of δ, K, C , as required. \square

With this fact in hand, we can now prove:

Proposition 5.13. *Let $\phi: \bar{E} \rightarrow \bar{E}$ be a quasi-isometry with $\nu_\phi^a = \text{id}_{E_0}$. There is a constant $C' = C'(\phi) > 0$ such that for all $x \in \bar{E}$, $d(x, \phi(x)) \leq C'$. Consequently, $\mathcal{A}: \text{QI}(\bar{E}) \rightarrow \text{Aff}(E_0)$ is injective.*

Proof. We first claim that for any $X \in \bar{D}$, $\phi(E_X)$ lies within the C'' -neighborhood of E_X , where $C'' = C''(\phi) > 0$. To see this, observe that since ν_ϕ^a is the identity and is bounded distance from ν_ϕ , it follows that $\phi|_{E_0}$ is within bounded distance of the inclusion of E_0 in \bar{E} . Proposition 5.4 then implies that $d_{\text{Haus}}(A, \phi_S(A)) < +\infty$ for each strip $A \in \mathcal{S}$. Since strips that lie within finite Hausdorff distance coincide, we have $\phi_S(A) = A$. Combining this fact with Proposition 5.4 it follows that $\phi_{\mathcal{P}}(\alpha(A)) = \alpha(\phi_S(A)) = \alpha(A)$. Hence for each α , we have that $\phi(\partial\mathcal{B}_\alpha)$ lies within Hausdorff distance C of $\partial\mathcal{B}_\alpha = \partial\mathcal{B}_{\phi_{\mathcal{P}}(\alpha)}$, for C as in Proposition 5.4.

Now let $X \in \bar{D}$ be any point and choose distinct $\alpha, \alpha' \in \mathcal{P}$ so that X is contained in the coarse intersection of $\partial\mathcal{B}_\alpha$ and $\partial\mathcal{B}_{\alpha'}$, implying that E_X lies in the coarse intersection of $\partial\mathcal{B}_\alpha$ and $\partial\mathcal{B}_{\alpha'}$. By the coarse preservation of the $\partial\mathcal{B}_\alpha$ in the previous paragraph, the coarse intersection of $\phi(\partial\mathcal{B}_\alpha)$ and $\phi(\partial\mathcal{B}_{\alpha'})$ is within Hausdorff distance C of the coarse intersection of $\partial\mathcal{B}_\alpha$ and $\partial\mathcal{B}_{\alpha'}$, and hence E_X and $\phi(E_X)$ are within uniform Hausdorff distance. This proves the claim.

Since $f_X: \bar{E} \rightarrow E_X$ is $e^{C''}$ -bi-Lipschitz when restricted to fibers in the C'' -neighborhood of E_X , the claim implies that $\nu_\phi^X = f_X \circ \phi: E_X \rightarrow E_X$ is a quasi-isometry with constants depending only on ϕ and not X . Moreover, since each E_X lies within finite (but not necessarily bounded) Hausdorff distance of E_0 , the fact that $\phi|_{E_0}$ lies within finite distance of the inclusion $E_0 \hookrightarrow \bar{E}$ implies that ν_ϕ^X lies within finite distance of the identity $E_X \rightarrow E_X$. Since each E_X is uniformly quasiisometric to \mathbb{H}^2 , it follows that $\nu_\phi^X: E_X \rightarrow E_X$ satisfies the assumptions of Lemma 5.12. We conclude that ν_ϕ^X is within uniformly bounded distance of the identity for each $X \in \bar{D}$. Since $d(\nu_\phi^X(x), \phi(x)) \leq C''$, it follows that $d(x, \phi(x))$ is uniformly bounded, independent of x . This proves the first statement of the proposition.

If $\mathcal{A}(\phi)$ is the identity for some $\phi \in \text{QI}(\bar{E})$, then by the first part of the proposition, ϕ is a bounded distance from the identity. Therefore, ϕ and the identity represent the same class, and \mathcal{A} is injective. This completes the proof. \square

5.5. From affine homeomorphisms to isometries. Next we will choose a particular allowable truncation and construct a homomorphism $\text{Aff}(E_0) \rightarrow \text{Isom}_{\text{fib}}(\bar{E})$, that we will eventually show is an isomorphism. We first construct such a homomorphism to the fiber-preserving isometry group of the space E , which avoids the issue of choosing the truncation.

Lemma 5.14. *For any $\nu \in \text{Aff}(E_0)$, there exists a isometry $\phi = \phi_\nu \in \text{Isom}_{\text{fib}}(E)$ such that $f_0 \circ \phi_\nu|_{E_0} = \nu$. Moreover, this assignment $\nu \mapsto \phi_\nu$ defines an injective homomorphism $\text{Aff}(E_0) \rightarrow \text{Isom}_{\text{fib}}(E)$.*

Proof. Recall from §2.1 that the projective tangent space at any non-cone point of E_0 is denoted $\mathbb{P}^1(q)$ and is canonically identified with ∂D . The derivative of $\nu: E_0 \rightarrow E_0$ (which may reverse orientations) is a well-defined projective transformation $d\nu \in \text{PGL}(\mathbb{P}^1(q))$ which, using the preferred coordinates on $q = q_0$ with distinguished vertical and horizontal directions, we canonically identify with $\text{PGL}_2(\mathbb{R})$. The Teichmüller disk D is the orbit of q under the $\text{SL}_2(\mathbb{R})$ action and is identified with $\mathbb{H}^2 = \text{SL}_2(\mathbb{R})/\text{SO}(2)$ (see e.g. [DDLS21, §2.8]). As $\text{PGL}_2(\mathbb{R})$ acts isometrically on \mathbb{H}^2 , we thus obtain an isometry $\Phi = d\nu: D \rightarrow D$ whose induced map $\partial\Phi$ of the circle at infinity ∂D agrees with the derivative $d\nu$ under the canonical identification $\partial D \cong \mathbb{P}^1(q)$. In particular, setting $X = \Phi(X_0)$, the geodesic ray in D emanating

from X_0 and asymptotic to $\xi \in \mathbb{P}^1(q)$ is sent to the geodesic ray emanating from X asymptotic to $d\nu(\xi)$.

We claim the map $\phi_0 = f_{X, X_0} \circ \nu: E_0 \rightarrow E_X$ is an isometry of fibers. Indeed, any pair $\xi, \xi^\perp \in \mathbb{P}^1(q)$ of orthogonal directions on E_0 are the endpoints of a geodesic ρ in D containing X_0 . Since Φ is an isometry with $\partial\Phi = d\nu$, we have that $X = \Phi(X_0)$ lies on the geodesic $\Phi(\rho)$ from $d\nu(\xi^\perp)$ to $d\nu(\xi)$; that is, $d\nu(\xi), d\nu(\xi^\perp)$ are orthogonal on X . But since $\mathbb{P}^1(q)$ and $\mathbb{P}^1(q_X)$ are canonically identified by the Teichmüller map f_{X, X_0} (see e.g. [DDLS21, §2.8]), this means ϕ_0 is an affine map whose derivative $d\phi_0 = d\nu$ preserves orthogonality of lines; hence ϕ_0 is an isometry as claimed.

Now we define $\phi = \phi_\nu: E \rightarrow E$ by the formula:

$$\phi(x) = f_{\Phi(\pi(x)), X_0} \circ \nu \circ f_0(x).$$

In words, this maps the fiber over a point Y to the fiber over $\Phi(Y)$, and the horizontal disk D_x , for $x \in E_0$, to $D_{\nu(x)}$. The restriction $\phi|_{D_x}: D_x \rightarrow D_{\nu(x)}$ is an isometry since it covers Φ . To prove that ϕ is an isometry, it therefore suffices to show that $\phi|_{E_Y}: E_Y \rightarrow E_{\Phi(Y)}$ is an isometry for any $Y \in D$.

Fix any $Y \in D$. For $Y = X_0$, we have already seen that $\phi|_{E_0}$ is the isometry $\phi_0: E_0 \rightarrow E_X$. If $Y \neq X$, there exist unique orthogonal directions $\alpha, \alpha^\perp \in \mathbb{P}(q)$ and $t > 0$, so that X_0 and Y both lie on the geodesic from α^\perp to α in D and Y lies distance t from X_0 in the direction of α . This means that $f_{Y, X_0}: E_0 \rightarrow E_Y$ contracts in direction α by e^{-t} and stretches in direction α^\perp by e^t . The image $\Phi(Y)$ lies along the geodesic from $d\phi_0(\alpha^\perp)$ to $d\phi_0(\alpha)$ at distance t from $\Phi(X_0) = X$; therefore $f_{\Phi(Y), X}: E_X \rightarrow E_{\Phi(Y)}$ contracts by e^{-t} in direction $d\phi_0(\alpha)$ and stretches by e^t in direction $d\phi_0(\alpha^\perp)$. The restriction $\phi|_Y: Y \rightarrow \Phi(Y)$ is given by $f_{\Phi(Y), X} \circ \phi_0 \circ f_{E_0, Y}$. Since ϕ_0 sends $\alpha \mapsto d\phi_0(\alpha)$ and $\alpha^\perp \mapsto d\phi_0(\alpha^\perp)$, the description above shows that $\phi|_Y$ is an isometry. Therefore ϕ is an isometry, as required.

To see that $\nu \mapsto \phi_\nu$ is a homomorphism, note that by construction Φ_ν is the unique isometry of D for which $\partial\Phi_\nu = d\nu$. Thus the chain rule implies $\Phi_{\nu \circ g} = \Phi_\nu \circ \Phi_g$ is the unique isometry whose action on ∂D agrees with $d(\nu \circ g) = d\nu \circ dg$. For any $x \in E$ we have $\pi(\phi_g(x)) = \Phi_g(\pi(x))$ and hence by construction

$$\begin{aligned} \phi_\nu \circ \phi_g(x) &= f_{\Phi_\nu(\pi(\phi_g(x))), X_0} \circ \nu \circ f_0(f_{\Phi_g(x), X_0} \circ g \circ f_0(x)) \\ &= f_{\Phi_\nu(\Phi_g(\pi(x))), X_0} \circ \nu \circ f_{X_0, \Phi_g(x)} \circ f_{\Phi_g(x), X_0} \circ g \circ f_0(x) \\ &= f_{\Phi_{\nu \circ g}(\pi(x)), X_0} \circ (\nu \circ g) \circ f_0(x) = \phi_{\nu \circ g}(x) \end{aligned}$$

as needed. Finally, if $\phi_\nu = \text{id}_E$ then clearly $X = \Phi(X_0) = X_0$. Since $\phi_0 = \phi_\nu|_{E_0}$ by construction, we conclude that

$$\text{id}_{E_0} = \phi_\nu|_{E_0} = \phi_0 = f_{X_0, X_0} \circ \nu = \nu.$$

Hence ν is the identity affine map, showing that $\nu \mapsto \phi_\nu$ is injective. \square

Lemma 5.15. *The subgroup $\Gamma < \text{Isom}_{\text{fib}}(E)$ has finite index.*

Proof. By [DDLS21, Proposition 5.5], $\text{Isom}_{\text{fib}}(E)$ acts properly discontinuously on E . Therefore $E/\text{Isom}_{\text{fib}}(E)$ is a topological orbifold with well-defined, positive Riemannian volume. The index of Γ in $\text{Isom}_{\text{fib}}(E)$ is the degree of the orbifold cover $E/\Gamma \rightarrow E/\text{Isom}_{\text{fib}}(E)$ and equals the ratio of the respective volumes. As E/Γ has finite volume, since the quotient D/G has finite area and the fibers $E_X/\pi_1 S$ all have equal, finite area, we conclude that Γ indeed has finite index. \square

Lemma 5.16. *There is an allowable truncation \bar{E} that is $\text{Isom}_{\text{fib}}(E)$ -invariant and for which restricting to \bar{E} induces an injection $\text{Isom}_{\text{fib}}(E) \rightarrow \text{Isom}_{\text{fib}}(\bar{E})$.*

Remark 5.17. Every fiber-preserving isometry of \bar{E} uniquely extends to one of E (e.g. by following the proof of Lemma 5.14) and thus $\text{Isom}_{\text{fib}}(E) \rightarrow \text{Isom}_{\text{fib}}(\bar{E})$ is in fact an isomorphism.

Proof. There is a natural map $\text{Isom}_{\text{fib}}(E) \rightarrow \text{Isom}(D)$ that sends Γ onto G . Hence, by the previous lemma, the image G^* of $\text{Isom}_{\text{fib}}(E)$ under this map contains G with finite index. Therefore G^* acts properly discontinuously on D and we may choose a collection $\{B_\alpha\}_{\alpha \in \mathcal{P}}$ of 1-separated horoballs as in §2.1 that is G^* -invariant. If \bar{E} denotes the corresponding truncation of E , it follows that every element of $\text{Isom}_{\text{fib}}(E)$ preserves \bar{E} . The map $\text{Isom}_{\text{fib}}(E) \rightarrow \text{Isom}_{\text{fib}}(\bar{E})$ given by restricting $\phi \mapsto \phi|_{\bar{E}}$ is injective by [DDLS21, Corollary 5.6] since if $\phi|_{\bar{E}}$ is the identity, then ϕ must be the identity on each Teichmüller disk D_x and fiber $E_X \subset \bar{E}$. \square

Choosing such an allowable truncation \bar{E} , Lemmas 5.14 and 5.16 now give an injective homomorphism $\Psi: \text{Aff}(E_0) \rightarrow \text{Isom}_{\text{fib}}(\bar{E})$ given by $\Psi(\nu) = \phi_\nu|_{\bar{E}}$.

Lemma 5.18. *For any $\nu \in \text{Aff}(E_0)$, $\mathcal{A}(\Psi(\nu)) = \nu$, where we have identified $\Psi(\nu)$ with its image in $\text{QI}(\bar{E})$ from the homomorphism $\text{Isom}_{\text{fib}}(\bar{E}) \rightarrow \text{QI}(\bar{E})$.*

Proof. The construction of \mathcal{A} in Corollary 5.9 sends the (quasi-)isometry $\Psi(\nu) = \phi_\nu|_{\bar{E}}: \bar{E} \rightarrow \bar{E}$ to the unique affine homeomorphism of E_0 that is uniformly close to the map $f_0 \circ \phi_\nu|_{E_0}: E_0 \rightarrow E_0$. But by the construction of ϕ_ν in Lemma 5.14, $f_0 \circ \phi_\nu|_{E_0}$ is affine itself and equal to ν . Thus evidently $\mathcal{A}(\Psi(\nu)) = \nu$ as claimed. \square

Lemma 5.19. *For any $\phi \in \text{Isom}_{\text{fib}}(\bar{E}) = \text{Isom}_{\text{fib}}(E)$, we have $\Psi(\mathcal{A}(\phi)) = \phi$. In particular, the natural maps $\text{Isom}_{\text{fib}}(\bar{E}) \rightarrow \text{Isom}(\bar{E}) \rightarrow \text{QI}(\bar{E})$ are both injective.*

Proof. By construction $\nu = \nu_\phi^a = \mathcal{A}(\phi)$ is the unique affine homeomorphism bounded distance from $f_0 \circ \phi|_{E_0}$. As this map is itself affine, we have $\nu = f_0 \circ \phi|_{E_0}$. The isometry $\Phi: D \rightarrow D$ in the construction of $\Psi(\nu)$ is then just the descent of ϕ to D . Further, for any $X, Y \in D$ we have $\phi|_X \circ f_{X,Y} = f_{\Phi(X),\Phi(Y)} \circ \phi|_{E_Y}$, since if X lies at distance $t > 0$ from Y along the geodesic from α^\perp to α , then both maps send $(\alpha^\perp, \alpha) \mapsto (\partial\Phi(\alpha^\perp), \partial\Phi(\alpha))$ while contracting the first by e^{-t} and expanding the second by e^t , hence they are the same affine map $E_Y \rightarrow E_{\Phi(X)}$. It follows that the restriction $\Psi(\nu)|_{E_Y}: E_Y \rightarrow E_{\Phi(Y)}$ is then the composition

$$\begin{aligned} \Psi(\nu)|_{E_Y} &= f_{\Phi(Y),X_0} \circ (f_0 \circ \phi|_{E_0}) \circ f_0|_{E_Y} = f_{\Phi(Y),\Phi(X_0)} \circ \phi|_{E_0} \circ f_{X_0,Y} \\ &= f_{\Phi(Y),\Phi(X_0)} \circ f_{\Phi(X_0),\Phi(Y)} \circ \phi|_{E_Y} = \phi|_{E_Y}. \end{aligned}$$

Since this holds for each Y , we conclude $\Psi(\nu) = \phi$ as claimed.

It follows that $\text{Isom}_{\text{fib}}(\bar{E}) \rightarrow \text{QI}(\bar{E})$ is injective, since if $[\phi]$ is the identity in $\text{QI}(\bar{E})$, meaning ϕ is finite distance from the identity, then $\mathcal{A}([\phi]) = \mathcal{A}(\text{id})$ and consequently $\phi = \Psi(\mathcal{A}([\phi]))$ are both the identity. Finally, $\text{Isom}(\bar{E}) \rightarrow \text{QI}(\bar{E})$ is injective since we have $\text{Isom}_{\text{fib}}(\bar{E}) = \text{Isom}(\bar{E})$ by [DDLS21, Corollary 5.4]. \square

5.6. Rigidity. We are now ready to complete the proof of Theorem 1.7:

Proof of Theorem 1.7. By Lemma 5.18, the composition

$$\text{Aff}(E_0) \xrightarrow{\Psi} \text{Isom}_{\text{fib}}(\bar{E}) \rightarrow \text{Isom}(\bar{E}) \rightarrow \text{QI}(\bar{E}) \cong \text{QI}(\Gamma) \xrightarrow{\mathcal{A}} \text{Aff}(E_0)$$

is the identity. Hence the first map Ψ is injective, and the remaining maps are injective by Lemma 5.19 and Proposition 5.13. It follows that each map above is an isomorphism, as claimed. The fact that Γ has finite index in $\text{Isom}(\bar{E}) \cong \text{QI}(\Gamma)$ thus follows from Lemma 5.15. \square

Standard techniques (see, for example [Sch95, §10.4]) now imply the following:

Corollary 5.20. *If H is any finitely generated group quasi-isometric to Γ , then H and Γ are weakly commensurable, meaning H has a finite normal subgroup N so that H/N and Γ contain finite index subgroups that are isomorphic.*

This proof requires one more lemma.

Lemma 5.21. *For every K there exists R' such that if $\phi: \bar{E} \rightarrow \bar{E}$ is a (K, K) -quasi-isometry that lies within finite distance of the identity, then ϕ lies within distance R' of the identity, meaning $d(x, \phi(x)) \leq R'$ for all $x \in \bar{E}$.*

Proof. First define a map $\bar{\phi}: \bar{D} \rightarrow \bar{D}$ by setting $\bar{\phi}(X) = Y$, where Y is the point provided by Proposition 5.4 such that $d_{\text{Haus}}(\phi(E_X), E_Y) \leq C$. Since $d_{\bar{D}}(X, Y) = d_{\bar{E}}(E_X, E_Y)$ for all X, Y in \bar{D} , we see that $\bar{\phi}$ is a quasi-isometry with constants depending only on K . It also lies within finite distance of the identity, as it inherits this property from ϕ ; thus applying Lemma 5.12 to $Z = \bar{D}$ implies that $\bar{\phi}$ lies within uniformly finite distance of the identity. That is, there exists R depending only on K so that $d_{\text{Haus}}(E_X, \bar{\phi}(E_X)) \leq R$ for all $X \in \bar{D}$. Hence, Lemma 5.6 implies that for each map $\nu_\phi^X = f_X \circ \phi|_{E_X}$ is a (K', K') -quasi-isometry for some K' depending only on K . Again by Lemma 5.12, this time with $Z = E_X$, we see that each ν_ϕ^X moves points uniformly bounded distance, and therefore $\phi|_{E_X}$ lies within uniform distance of the inclusion of E_X in \bar{E} . Since this holds for all X , we have that ϕ lies within uniform distance of the identity, as required. \square

Proof of Corollary 5.20. If H is quasi-isometric to Γ , there is a quasi-isometry $\mu: H \rightarrow \bar{E}$ with a quasi-inverse $\mu^{-1}: \bar{E} \rightarrow H$. Left multiplication by $h \in H$ gives an isometry $L_h: H \rightarrow H$. In this way, for each $h \in H$ we obtain a quasi-isometry $\mathcal{B}(h) = \mu \circ L_h \circ \mu^{-1}$ of \bar{E} with *uniformly bounded constants*. Let us also set $\mathcal{B}'(h) = \Psi(\mathcal{A}(\mathcal{B}(h))) \in \text{Isom}_{\text{fib}}(\bar{E}) = \text{Isom}(\bar{E})$, which is the unique isometry of \bar{E} at finite distance from $\mathcal{B}(h)$. Since the quasi-isometry constants of $\mathcal{B}(h)$ are uniform, depending only on μ , it follows from Lemma 5.21 that there is a constant R' so that $d(\mathcal{B}(h)(x), \mathcal{B}'(h)(x)) \leq R'$ for all $x \in \bar{E}$ and $h \in H$.

We now claim the homomorphism $\mathcal{B}': H \rightarrow \text{Isom}(\bar{E})$ has finite kernel and cokernel. Indeed, if $\mathcal{B}'(h) = \text{Id}_{\bar{E}}$ the above implies $\mathcal{B}(h)$ moves $\mu(e)$ (and in fact all points) distance at most R' . But this means L_h moves the identity $e \in H$ uniformly bounded distance, and there are only finitely many such elements of H . To prove \mathcal{B}' has finite cokernel it suffices, as in Lemma 5.15, to show $\bar{E}/\mathcal{B}'(H)$ has finite volume or, better yet, finite diameter. For this, given $x, y \in \bar{E}$ we must find h so that $d(\mathcal{B}'(h)(x), y)$ is uniformly bounded. This is equivalent to bounding $d(\mathcal{B}(h)(x), y) = d(\mu(h \cdot \mu^{-1}(x)), y)$, which is coarsely $d_H(h \cdot \mu^{-1}(x), \mu^{-1}(y))$. Since H acts transitively on itself, this is clearly possible.

We now see that $H/\ker(\mathcal{B}')$ and Γ are both realized as finite index subgroups of $\text{Isom}(\bar{E})$ and hence that their intersection has finite index in both. \square

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*, 8:66–104, 2021.
- [ABO19] Carolyn Abbott, Sahana H. Balasubramanya, and Denis Osin. Hyperbolic structures on groups. *Algebr. Geom. Topol.*, 19(4):1747–1835, 2019.
- [Ahl66] Lars V. Ahlfors. Fundamental polyhedrons and limit point sets of Kleinian groups. *Proc. Nat. Acad. Sci. U.S.A.*, 55:251–254, 1966.
- [BBKL20] Mladen Bestvina, Kenneth Bromberg, Autumn E. Kent, and Christopher J. Leininger. Undistorted purely pseudo-Anosov groups. *J. Reine Angew. Math.*, 760:213–227, 2020.
- [BHMS20] Jason Behrstock, Mark F. Hagen, A. Martin, and A. Sisto. A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups. preprint, [arXiv:2005.00567](https://arxiv.org/abs/2005.00567), 2020.
- [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(5):890–926, 2017.
- [BHS17b] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.
- [BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299(2):257–338, 2019.
- [BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.*, 170(5):–, 2021.
- [BL18] Anja Bankovic and Christopher J. Leininger. Marked-length-spectral rigidity for flat metrics. *Trans. Amer. Math. Soc.*, 370(3):1867–1884, 2018.
- [Bow93] B. H. Bowditch. Geometrical finiteness for hyperbolic groups. *J. Funct. Anal.*, 113(2):245–317, 1993.
- [Bow13] Brian H. Bowditch. Coarse median spaces and groups. *Pacific J. Math.*, 261(1):53–93, 2013.
- [Bow18] Brian H. Bowditch. Large-scale rigidity properties of the mapping class groups. *Pacific J. Math.*, 293(1):1–73, 2018.
- [CLM12] Matt T. Clay, Christopher J. Leininger, and Johanna Mangahas. The geometry of right-angled Artin subgroups of mapping class groups. *Groups Geom. Dyn.*, 6(2):249–278, 2012.
- [DDL21] Spencer Dowdall, Matthew G. Durham, Christopher J. Leininger, and Alessandro Sisto. Extensions of Veech groups I: A hyperbolic action. *arXiv preprint arXiv:2006.16425v2*, 2021.
- [DELS18] Moon Duchin, Viveka Erlandsson, Christopher J. Leininger, and Chandrika Sadanand. You can hear the shape of a billiard table: Symbolic dynamics and rigidity for flat surfaces. Preprint, to appear in *Comment. Math. Helv.*, 2018.
- [DGO17] F. Dahmani, V. Guirardel, and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.*, 245(1156):v+152, 2017.
- [DHS17] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.
- [DMS20] Matthew G Durham, Yair N Minsky, and Alessandro Sisto. Stable cubulations, bicomings, and barycenters. *arXiv preprint arXiv:2009.13647*, 2020.
- [DT15] Matthew Gentry Durham and Samuel J. Taylor. Convex cocompactness and stability in mapping class groups. *Algebr. Geom. Topol.*, 15(5):2839–2859, 2015.
- [EF97] David B. A. Epstein and Koji Fujiwara. The second bounded cohomology of word-hyperbolic groups. *Topology*, 36(6):1275–1289, 1997.
- [Far98] B. Farb. Relatively hyperbolic groups. *Geom. Funct. Anal.*, 8(5):810–840, 1998.
- [FM02a] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152, 2002.
- [FM02b] Benson Farb and Lee Mosher. The geometry of surface-by-free groups. *Geometric & Functional Analysis GAFA*, 12(5):915–963, 2002.

- [Gre66] L. Greenberg. Fundamental polyhedra for kleinian groups. *Ann. of Math. (2)*, 84:433–441, 1966.
- [Ham] Ursula Hamenstädt. Word hyperbolic extensions of surface groups. Preprint, [arXiv:math.GT/0505244](https://arxiv.org/abs/math/0505244).
- [HHP20] Thomas Haettel, Nima Hoda, and Harry Petyt. Coarse injectivity, hierarchical hyperbolicity, and semihyperbolicity. *arXiv preprint arXiv:2009.14053*, 2020. To appear in *Geom. Topol.*
- [HO13] Michael Hull and Denis Osin. Induced quasicocycles on groups with hyperbolically embedded subgroups. *Algebr. Geom. Topol.*, 13(5):2635–2665, 2013.
- [HRSS21] Mark F. Hagen, Jacob Russell, Alessandro Sisto, and Davide Spriano. Equivariant hierarchically hyperbolic structures for 3-manifold groups via quasimorphisms. *In preparation*, 2021.
- [JS79] William H. Jaco and Peter B. Shalen. Seifert fibered spaces in 3-manifolds. *Mem. Amer. Math. Soc.*, 21(220):viii+192, 1979.
- [KL97] Michael Kapovich and Bernhard Leeb. Quasi-isometries preserve the geometric decomposition of Haken manifolds. *Invent. Math.*, 128(2):393–416, 1997.
- [KL07] Autumn E. Kent and Christopher J. Leininger. Subgroups of mapping class groups from the geometrical viewpoint. In *In the tradition of Ahlfors-Bers. IV*, volume 432 of *Contemp. Math.*, pages 119–141. Amer. Math. Soc., Providence, RI, 2007.
- [KL08a] Autumn E. Kent and Christopher J. Leininger. Shadows of mapping class groups: capturing convex cocompactness. *Geom. Funct. Anal.*, 18(4):1270–1325, 2008.
- [KL08b] Autumn E. Kent and Christopher J. Leininger. Uniform convergence in the mapping class group. *Ergodic Theory Dynam. Systems*, 28(4):1177–1195, 2008.
- [KL12] Ilya Kapovich and Anton Lukyanenko. Quasi-isometric co-Hopficity of non-uniform lattices in rank-one semi-simple Lie groups. *Conform. Geom. Dyn.*, 16:269–282, 2012.
- [Kob12] Thomas Koberda. Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups. *Geom. Funct. Anal.*, 22(6):1541–1590, 2012.
- [Loa21] Christopher Loa. Free products of abelian groups in mapping class groups. *arXiv preprint arXiv:2103.05144*, 2021.
- [LR06] C. J. Leininger and A. W. Reid. A combination theorem for Veech subgroups of the mapping class group. *Geom. Funct. Anal.*, 16(2):403–436, 2006.
- [Man05] Jason Fox Manning. Geometry of pseudocharacters. *Geom. Topol.*, 9:1147–1185, 2005.
- [Mar74] Albert Marden. The geometry of finitely generated kleinian groups. *Ann. of Math. (2)*, 99:383–462, 1974.
- [Mas70] Bernard Maskit. On boundaries of Teichmüller spaces and on Kleinian groups. II. *Ann. of Math. (2)*, 91:607–639, 1970.
- [Min96a] Yair N. Minsky. Extremal length estimates and product regions in Teichmüller space. *Duke Math. J.*, 83(2):249–286, 1996.
- [Min96b] Yair N. Minsky. Quasi-projections in Teichmüller space. *J. Reine Angew. Math.*, 473:121–136, 1996.
- [MM99] H.A. Masur and Y.N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999.
- [MM00] H.A. Masur and Y. N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.
- [Mos06] Lee Mosher. Problems in the geometry of surface group extensions. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 245–256. Amer. Math. Soc., Providence, RI, 2006.
- [MS12] Mahan Mj and Pranab Sardar. A combination theorem for metric bundles. *Geom. Funct. Anal.*, 22(6):1636–1707, 2012.
- [MT02] Howard Masur and Serge Tabachnikov. Rational billiards and flat structures. In *Handbook of dynamical systems, Vol. 1A*, pages 1015–1089. North-Holland, Amsterdam, 2002.
- [MW95] H.A. Masur and M. Wolf. Teichmüller space is not Gromov hyperbolic. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 20(2):259–267, 1995.
- [Raf14] Kasra Rafi. Hyperbolicity in Teichmüller space. *Geom. Topol.*, 18(5):3025–3053, 2014.

- [RST18] Jacob Russell, Davide Spriano, and Hung Cong Tran. Convexity in hierarchically hyperbolic spaces. *arXiv preprint arXiv:1809.09303*, 2018. To appear in *Algebr. Geom. Topol.*
- [Run20] Ian Runnels. Effective generation of right-angled artin groups in mapping class groups. *arXiv preprint arXiv:2004.13585*, 2020.
- [Rus21] Jacob Russell. Extensions of multicurve stabilizers are hierarchically hyperbolic. *arXiv preprint arXiv:2107.14116*, 2021.
- [Sch95] Richard Evan Schwartz. The quasi-isometry classification of rank one lattices. *Inst. Hautes Études Sci. Publ. Math.*, (82):133–168 (1996), 1995.
- [Sis16] Alessandro Sisto. Quasi-convexity of hyperbolically embedded subgroups. *Math. Z.*, 283(3-4):649–658, 2016.
- [Sis19] Alessandro Sisto. What is a hierarchically hyperbolic space? In *Beyond hyperbolicity*, volume 454 of *London Math. Soc. Lecture Note Ser.*, pages 117–148. Cambridge Univ. Press, Cambridge, 2019.
- [Tan19] Robert Tang. Affine diffeomorphism groups are undistorted. *arXiv preprint arXiv:1912.06537*, 2019.
- [Thu86] W.P. Thurston. Geometry and topology of 3-manifolds. Princeton University Lecture Notes, online at <http://www.msri.org/publications/books/gt3m>, 1986.

DEPARTMENT OF MATHEMATICS, VANDERBILT UNIVERSITY, NASHVILLE, TN
Email address: `spencer.dowdall@vanderbilt.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, RIVERSIDE, CA
Email address: `mdurham@ucr.edu`

DEPARTMENT OF MATHEMATICS, RICE UNIVERSITY, HOUSTON, TX
Email address: `cjl12@rice.edu`

DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY, EDINBURGH, UK
Email address: `a.sisto@hw.ac.uk`