Extensions of Multicurve Stabilizers are Hierarchically Hyperbolic

Jacob Russell

19th September, 2023

Abstract

For a closed and orientable surface S with genus at least 2, we prove the $\pi_1(S)$-extensions of the stabilizers of multicurves on S are hierarchically hyperbolic groups. This answers a question of Durham, Dowdall, Leininger, and Sisto. We also include an appendix that employs work of Charney, Cordes, and Sisto to characterize the Morse boundaries of hierarchically hyperbolic groups whose largest acylindrical action on a hyperbolic space is on a quasi-tree.

Contents

1 Introduction

For an orientable closed surface S and any group G, a $\pi_1(S)$-extension of G is any group E that fits into the short exact sequence

$$1 \to \pi_1(S) \to E \to G \to 1.$$

Topologically, $\pi_1(S)$-extensions arise as the fundamental groups of surface bundles.

When S has genus at least 2, the Birman exact sequence tells us that the mapping class group of S and the mapping class group fixing a marked point $z \in S$ fit into the short exact sequence

$$1 \to \pi_1(S) \to \mathrm{MCG}(S; z) \to \mathrm{MCG}(S) \to 1$$

where $\pi_1(S)$ is identified with the point pushing subgroup of $\mathrm{MCG}(S; z)$. This allows us to create a $\pi_1(S)$-extension of any subgroup $G < \mathrm{MCG}(S)$ by taking the full preimage of G in $\mathrm{MCG}(S; z)$.

Because the mapping class group of S is naturally isomorphic to the outer automorphism group of $\pi_1(S)$, every $\pi_1(S)$-extension of an abstract group G produces a monodromy homomorphism $G \to \mathrm{MCG}(S)$. The extension groups arising from the Birman exact sequence are therefore the extensions whose monodromies have trivial kernel. Since a $\pi_1(S)$-extension is determined up to isomorphism by the monodromy into $\mathrm{MCG}(S)$, understanding how properties of the monodromy influence the properties of the $\pi_1(S)$-extension group is an important problem in mapping class groups.

The first examples of $\pi_1(S)$-extensions arise from 3-manifolds that fiber over the circle with fiber S. In these examples, $G \cong Z$ and, by Thurston's Geometerization Theorem, any such 3-manifold admits a hyperbolic structure if and only if the image of the monodromy is generated by a pseudo-Anosov element of $\mathrm{MCG}(S)$; see [Ota96]. Work of Farb and Mosher [FM02] plus Hamenstädt [Ham] expanded Thurston's result to prove that a $\pi_1(S)$-extension of a group G is *Gromov hyperbolic* if and only if the monodromy has finite kernel and *convex cocompact* image; see also [MS12].

In the present work, we study the geometry of $\pi_1(S)$-extensions when G is the stabilizer in $\mathrm{MCG}(S)$ of a multicurve on S. While the presence of Dehn twists in G prevents these extensions from ever being Gromov hyperbolic, we prove they are *hierarchically hyperbolic*. This gives a positive answer to a question of Dowdall, Durham, Leininger, and Sisto [DDLS, Question 1.12] asked as part of a search for a robust definition of geometric finiteness in the mapping class group; see Section 1.1 for further discussion of this motivation.

Theorem 1.1. *Let S be a closed orientable surface with genus at least 2. Let α be a multicurve on S and G_α be the stabilizer of α in $\mathrm{MCG}(S)$. If E_α is the full preimage of G_α in $\mathrm{MCG}(S; z)$, then E_α is a hierarchically hyperbolic group.*

Hierarchical hyperbolicity was introduced by Behrstock, Hagen, and Sisto to axiomatize the coarse geometric structure of the mapping class group arising from the machinery of Masur and Minsky [BHS17b, BHS19]. Masur and Minsky proved that the curve complex is Gromov hyperbolic and used projections of $\mathrm{MCG}(S)$ onto the curve complexes of subsurfaces of S to greatly illuminate the geometry of the mapping class group [MM99, MM00]. A number of subsequent results built on this work, producing a beautify theory of how the geometry of $\mathrm{MCG}(S)$ can be decoded from a combination of the geometry of these curve complexes and combinatorial information about the subsurfaces of S, e.g., [Beh06, BDM09, BKMM12, BM08, Bow14]. Hierarchical hyperbolicity axiomatizes this theory, describing a class of spaces whose coarse geometry is encoded in a collection of projections onto hyperbolic metric spaces that are organized by a set of combinatorial relations. Remarkably, the class of hierarchically hyperbolic spaces encompasses a variety of groups beyond the mapping class group including the fundamental group of most 3-manifolds [BHS19], many cocompactly cubulated groups [BHS17b, HS20], Artin groups of extra large type [HMS], and several combinations of hyperbolic groups [BR20, RS, **?**]. Hierarchical hyperbolicity also describes the coarse geometry of a number of other groups and spaces associated to surfaces such as Teichmüller space with both the Teichmüller and Weil–Peterson metrics [BHS17b, MM99, MM00, BKMM12, Bro03, Dur16, Raf07, EMR17], the genus 2 handlebody group [Che22], the $\pi_1(S)$-extensions of lattice Veech groups [DDLS], certain quotients of the mapping class group [BHS17a, BHMS20], and a wide variety of graphs built from curves on surfaces [Vok22].

Hierarchical hyperbolicity produces a large number of geometric and algebraic consequences, e.g., [BHS17b, BHS19, BHS21, HHP20, RST23]. The following corollary states some salient examples of the new results that are gained automatically as a result of the hierarchical hyperbolicity of E_α.

Corollary 1.2. *Let S be a closed orientable surface with genus at least 2. Let α be a multicurve on S and let E_α be the full preimage in $\mathrm{MCG}(S; z)$ of the stabilizer of α in $\mathrm{MCG}(S)$.*

1. *E_α has quadratic Dehn function.*
2. *E_α has finitely many conjugacy classes of finite order subgroups.*
3. *E_α is semi-hyperbolic and hence has solvable conjugacy problem.*

Our proof of Theorem 1.1 uses the recent *combinatorial hierarchical hyperbolicity* machinery of Behrstock, Hagen, Martin, and Sisto [BHMS20]. The key to this approach is to

construct a hyperbolic simplicial complex X_α for E_α that is analogous to the curve complex for $\mathrm{MCG}(S)$. To define X_α, let S^z denote the surface obtained from S by adding $z \in S$ as a marked point, and let $\Pi\colon S^z \to S$ be the map given by forgetting that z is a marked point. The vertices of X_α are all isotopy classes of curves c on S^z so that $\Pi(c)$ is either contained in or disjoint from α. There are edges between two vertices of X_α if the isotopy classes have disjoint representatives.

We prove that X_α is not only hyperbolic, but in fact a quasi-tree. Combining this with results on hierarchical hyperbolicity from the literature produces several additional properties of E_α.

Theorem 1.3. *The graph X_α is uniformly quasi-isometric to a tree and the group E_α has the following properties.*

- *E_α acts acylindrically on X_α and this action is largest in the sense of* [ABO19].
- *A subgroup $H < E_\alpha$ is stable if and only if the orbit map of H in X_α is a quasi-isometric embedding. In particular, every stable subgroup of E_α is virtually free.*
- *The Morse boundary of E_α is an ω-Cantor space.*

The proof that the Morse boundary of E_α is an ω-Cantor space uses a technique developed by Charney, Cordes, and Sisto [CCS23]. In the appendix, we work out this application not just for E_α, but for all hierarchically hyperbolic groups whose largest acylindrical action on a hyperbolic space is on a quasi-tree.

1.1 Geometric finiteness in $\mathrm{MCG}(S)$

There is a longstanding and fruitful analogy between discrete subgroups of $\mathrm{Isom}(\mathbb{H}^n)$ and the subgroups of $\mathrm{MCG}(S)$; see [KL07] for an detailed overview. In the case of $\mathrm{Isom}(\mathbb{H}^n)$, the best behaved discrete subgroups are the geometrically finite subgroups and the convex cocompact subgroups. A discrete subgroup $\Gamma < \mathrm{Isom}(\mathbb{H}^n)$ is *geometrically finite* if Γ acts with finite co-volume on the convex hull of its limit set, and it is *convex cocompact* if it instead acts cocompactly.

Building on this analogy, Farb and Mosher defined a *convex cocompact* subgroup of $\mathrm{MCG}(S)$ as a subgroup whose orbit in the Teichmüller space of S is quasiconvex [FM02]. Convex cocompact subgroups have since seen immense study, yielding a variety of characterizations [KL08a, KL08b, DT15, BBKL20]. Notable to the study of $\pi_1(S)$-extensions, a subgroup of $\mathrm{MCG}(S)$ is convex cocompact if and only if the full preimage in $\mathrm{MCG}(S; z)$ is Gromov hyperbolic [FM02, Ham, MS12].

While convex cocompactness has been well studied in the mapping class group, the lack of negative curvature in the Teichmüller space of S has prevented the formulation of a robust analogue of geometric finiteness for $\mathrm{MCG}(S)$. Despite this, there are several subgroups of $\mathrm{MCG}(S)$ that one could naturally consider as geometrically finite. The most notable examples are the Veech groups, all of $\mathrm{MCG}(S)$ itself, and the stabilizers of multicurves on S. These subgroup are all not convex cocompact, but each acts with finite co-volume on a well-behaved subset of the Teichmüller space of S.

Dowdall, Durham, Leininger, and Sisto proposed that geometric finiteness (however it is defined) in $\mathrm{MCG}(S)$ should be characterized by some sort of hyperbolicity of the extension group that encompasses Gromov hyperbolicity in the convex cocompact case [DDLS]. In support of this idea, they proved that the full preimage in $\mathrm{MCG}(S; z)$ of a lattice Veech group is hierarchically hyperbolic and asked if the same could be proven for other candidates of geometrically finite subgroups [DDLS, Question 1.12]. Theorem 1.1 affirmatively answers this question for the stabilizers of multicurves. Combining these results with the observation that the full preimage of the entire $\mathrm{MCG}(S)$ is $\mathrm{MCG}(S; z)$—which is hierarchically hyperbolic—we now know that all three of the most naturally geometrically finite subgroups have hierarchically hyperbolic extension groups.

1.2 Outline of the paper

Section 2 reviews background material and notation on simplicial complexes, coarse geometry, group actions on graphs, curves on surfaces, and combinatorial hierarchically hyperbolic spaces. In Section 3, we provide a sketch of the hierarchical hyperbolicity of E_α when α is a single curve. Section 4 presents the criteria we will use to prove the graph X_α is a quasi-tree.

The proof that E_α is a hierarchically hyperbolic group is spread over two sections. In Section 5, we construct the quasi-tree X_α and use it to build a combinatorial HHS that E_α acts on non-properly. The source of this non-properness is the fact that Dehn twists in E_α will stabilize simplices in X_α. In Section 6, we address this lack of properness by constructing a new combinatorial HHS using a "blow-up" of the graph X_α. This blow-up of X_α will prevent Dehn twists from fixing simplices by recording the action of twists around curves in X_α. This in-turn causes E_α to act metrically properly on the resulting combinatorial HHS.

Section 7 is largely expository, elaborating on the hierarchically hyperbolic structure of E_α and some of its consequences. We start by giving a description of the hierarchically hyperbolic group structure on E_α using the topology of curves and subsurfaces. We then use this description to show that a minor modification of the HHG structure produces an HHG structure for E_α that has an additional property introduced by Abbott, Behrstock, and Durham called *unbounded products*. This allows us to use results from the literature plus the fact that X_α is a quasi-tree to achieve the conclusions of Theorem 1.3.

We include an appendix that shows how a technique developed by Charney, Cordes, and Sisto to understand the Morse boundary of right-angled Artin groups and graph manifolds groups can be extended to certain hierarchically hyperbolic groups. This extension is required to conclude the Morse boundary of E_α is an ω-Cantor space.

Acknowledgments

The author is very grateful to Chris Leininger for abundant conversations on the topic and for detailed comments on drafts of this paper. The author also thanks Alessandro Sisto and Mark Hagen for useful conversations about combinatorial HHSs. The author is grateful for the work of an anonymous referee whose many comments improved the exposition of this paper.

2 Background and notation

2.1 Simplicial Complexes

Throughout, if X is a simplicial complex or graph, then X^0 will denote the set of vertices of X. If the notation for the simplicial complex or graph depends on some parameter, for example $\mathcal{C}(S)$, then we will instead insert the superscript 0 between the identifying symbol and the parameter, for example, $\mathcal{C}^0(S)$ is the vertices of the complex $\mathcal{C}(S)$. We equip each graph with the path metric coming from declaring each edge to have length 1. When considering a simplicial complex as a metric space we will implicitly be referring to the metric space that is the 1-skeleton of the simplicial complex equipped with this path metric.

We will make frequent use of the following notions of join, link, and star.

Definition 2.1 (Join, link, and star)**.** Let X be a flag simplicial complex. If Y and Z are disjoint flag subcomplexes of X so that every vertex of Y is joined by an edge to every vertex of Z, then the *join* of Y and Z, $Y \bowtie Z$, is the subcomplex of X spanned by Y and Z. Given a subcomplex Y of X, the *link* of Y, $\mathrm{lk}(Y)$, is the subcomplex of X spanned by the vertices of X that are joined by an edge to all the vertices of Y. The *star* of Y, $\mathrm{st}(Y)$, is the join $Y \bowtie \mathrm{lk}(Y)$.

We say a simplex Δ of a simplicial complex X is *maximal* if $\mathrm{lk}(\Delta) = \varnothing$. This is equivalent to saying $\Delta \subseteq \Delta'$ implies $\Delta = \Delta'$ for any simplex Δ'. Note that the maximal simplices of X need not all have the same number of vertices. Finally, if $Y \subseteq X^0$, then we let $X - Y$ denote the subcomplex of X spanned by the vertices of X that are not in Y.

2.2 Coarse geometry and hyperbolic spaces

Throughout, when we measure the distances between sets in a metric space (or between a set a point), we are taking the minimum distance between the two sets.

A map $f\colon X \to Y$ between two metric spaces is a (λ, ϵ)-*quasi-isometric embedding* if $\lambda \geqslant 1$ and $\epsilon \geqslant 0$ and for all $x_1, x_2 \in X$ we have

$$\frac{1}{\lambda} d_X(x_1, x_2) - \epsilon \leqslant d_Y(f(x_1), f(x_2)) \leqslant \lambda d_X(x_1, x_2) + \epsilon.$$

A (λ, ϵ)-quasi-isometric embedding $f\colon X \to Y$ is a (λ, ϵ)-*quasi-isometry* if for each $y \in Y$, there exists $x \in X$ with $d_Y(y, f(x)) \leqslant \epsilon$. Because we will largely be interested in geometry up to quasi-isometry, we will often work with coarse maps instead of genuine functions. A *coarse map* between metric spaces X and Y is a function $f\colon X \to 2^Y$ so that $f(x)$ is a non-empty subset of Y with uniformly bounded diameter for all $x \in X$. By abuse of notation, we will often use the notation $f\colon X \to Y$ to signify coarse maps. We say a coarse map is a (λ, ϵ)-quasi-isometric embedding (resp. a quasi-isometry) if it satisfies the same inequalities given above (with distances between sets being the minimum distance). A coarse map $f\colon X \to Y$ is (λ, ϵ)-*coarsely Lipschitz* if

$$d_Y(f(x_1), f(x_2)) \leqslant \lambda d_X(x_1, x_2) + \epsilon$$

for each $x_1, x_2 \in X$.

When X is a graph, a coarse map defined on the vertices X^0 induces a coarse map on the entire graph by sending points on edges to the union of the images of the vertices of that edge. In this case, we freely use the following lemma to ensure that the induced map is a coarsely Lipschitz coarse map.

Lemma 2.2. *Let X be a connected graph and $f\colon X^0 \to Y$ be a coarse map. If there exist $\lambda \geqslant 1$ so that $\operatorname{diam}_Y(f(x)) \leqslant \lambda$ for each x and $d_Y(f(x_1), f(x_2)) \leqslant \lambda$ for every $x_1, x_2 \in X^0$ that are joined by an edge, then the map $f\colon X \to Y$ induced by the map on the vertices is $(3\lambda, 0)$-coarsely Lipschitz coarse map.*

Proof. Extend f to points on edges by defining the image of an edge point to be the union of the images of the vertices of that edge. This is a coarse map by the hypotheses. Let $x, x' \in X$ and $x_0, x_1, \ldots, x_n$ be the vertices on a geodesic in X connecting x to x'. Since we are measuring distance between sets as the minimum distance, we have

$$\begin{aligned} d_Y(f(x), f(x')) \leqslant d_Y(f(x_0), f(x_n)) &\leqslant \sum_{i=1}^{n} \operatorname{diam}_Y(f(x_{i-1}) \cup f(x_i)) \\ &\leqslant 3\lambda n = 3\lambda d_X(x_0, x_n) \leqslant 3\lambda d_X(x, x'). \end{aligned}$$

□

We will also frequently use the following criteria to prove that subsets of a graph are quasi-isometrically embedded.

Lemma 2.3. *Let H and Y be connected graphs equipped with the path metric. Suppose that H is a subgraph of Y. If there exists a coarse map $\psi\colon Y \to H$ so that*

- *ψ is the identity on H^0;*
- *ψ is (λ, ϵ)-coarsely Lipschitz;*

then, the inclusion of H into Y is a $(\lambda, \epsilon + 2\lambda)$-quasi-isometric embedding.

Proof. Let $h_1, h_2 \in H^0$. Since H is a connected subgraph, we have $d_Y(h_1, h_2) \leqslant d_H(h_1, h_2)$. Since $\psi(h_1) = h_1$ and $\psi(h_2) = h_2$ we have $d_H(h_1, h_2) \leqslant \lambda d_Y(h_1, h_2) + \epsilon$, hence the inclusion is a (λ, ϵ)-quasi-isometry on the vertices. Since each point on an edge is at most 1 from two vertices, this implies the inclusion is a $(\lambda, \epsilon + 2\lambda)$-quasi-isometric embedding on all of H. □

We say that a geodesic metric space is *δ-hyperbolic* if for every geodesic triangle in the space, the δ-neighborhood of any two sides covers the third side. If X and Y are quasi-isometric geodesic metric spaces and X is δ-hyperbolic, then Y is δ'-hyperbolic for some δ' depending on δ and the constants of the quasi-isometry $X \to Y$. A connected graph X is 0-hyperbolic if and only if X is a tree. We say that a geodesic metric space is a *quasi-tree* if it is quasi-isometric to a tree.

If X is a δ-hyperbolic graph and H is a connected subgraph so that the inclusion $H \to X$ is a (λ, ϵ)-quasi-isometric embedding, then for each $x \in X$ the set

$$\{h \in H : d_X(x, h) = d_X(x, H)\}$$

is uniformly bounded in terms of δ, λ, and ϵ. We therefore have a coarse map $\mathfrak{p}\colon X \to H$ given by $\mathfrak{p}(x) = \{h \in H : d_X(x, h) = d_X(x, H)\}$ (note, $\mathfrak{p}$ is the identity on H). We call this map the *coarse closest point projection* onto H. The coarse closest point projection is always coarsely Lipschitz [BH99, Proposition III.Γ.3.11].

2.3 Groups acting on graphs

Let G be a finitely generated group acting on a connected simplicial graph X by graph automorphisms. Let $\mathrm{ball}_r(x)$ denote the closed ball of radius $r \geqslant 0$ in X around the vertex $x \in X^0$. The action of G on X is *metrically proper* if for every $r \geqslant 0$ and $x \in X^0$, the set

$$\{g \in G : g \cdot \mathrm{ball}_r(x) \cap \mathrm{ball}_r(x) \neq \varnothing\}$$

is finite. The action of G on X is *cocompact* if the quotient of X by the action of G is compact. Similarly, the action is *cobounded* if the quotient has finite diameter.

2.4 Surfaces, curves, and mapping class groups

Let $S^n_{g,p}$ denote a connected orientable surface with genus g, p punctures, and n marked points. The complexity of $S^n_{g,p}$ is $\xi(S^n_{g,p}) = 3g - 3 + p + n$. Given a non-marked point z on the surface $S \cong S^n_{g,p}$, let S^z denote the surface obtained from S by adding z as an additional mark point.

By a *curve on* S, we mean an isotopy class of a simple closed curve on S that is essential on S minus the marked points of S. Here, two curves are considered isotopic if they are isotopic on S minus the marked points. We say two curves are *disjoint* if they are not equal and have disjoint representatives. A *multicurve* on S is a set of pairwise disjoint curves on S. For two curves c and c', we define the *intersection number*, $i(c, c')$, to be the minimal number of intersection points between representatives of c and c'. Note, $i(c, c') = 0$ means either $c = c'$ or c and c' are disjoint.

By a *subsurface* of S, we mean an isotopy class of an open essential subsurface of S, where two subsurfaces are considered isotopic if they are isotopic on S minus the marked points. Subsurfaces of S are not required to be connected. For a subsurface $U \subseteq S$, we use ∂U to denote the isotopy class of the curves on S that are the boundary curves for any representative of U. If A is an annular subsurface of S, then the *core curve* of A is the isotopy class of ∂A.

The set of all subsurface of S has a partial order denoted $\subseteq$ and a difference denoted $-$. For connected non-annular subsurfaces U and V, we write $U \subseteq V$ if U has a representative that is contained in a representative of V. In this case, we can subtract U from V as follows: fix a complete hyperbolic metric on S minus the marked points and pick representatives for V and U whose boundary curves are geodesic. The subsurface $V - U$ is the open subsurface that is the interior of the difference of these representatives of V and U. For an annular subsurface A and a connected subsurface V, we write $A \subseteq V$ if either $V = A$ or V is not an annulus and there is a representative of A that is a non-peripheral annulus on a representative of V. If $V = A$, then $V - A = \varnothing$. If $A \subsetneq V$, then there is a representative of ∂A that is contained in a representative of V. Define $V - A$ to be the isotopy class of the subsurface obtained by subtracting this representative of ∂A from this representative of V. If U and V are possibly disconnected surfaces, then $U \subseteq V$ if every component of U is contained in some component of V, where containment is defined in the above case of connected subsurfaces. Similarly, $V - U$ is defined for disconnected surfaces with $U \subseteq V$ by extending the difference defined above over connected components.

There is also a containment relation between multicurves and subsurfaces as well as a difference operation whenever a subsurface contains a multicurve. If U is a connected non-annular subsurface and a curve c has a representative that is a non-peripheral curve on some representative of U, then we say c is *contained* in U and write $c \subseteq U$. If A is an annulus, then we declare that the only curve that is contained in A is the core curve of A and we write $c \subseteq A$ only if c is the core curve of A. If every component of a multicurve μ is contained in a (possibly disconnected) subsurface U, then we say μ is *contained* in U and write $\mu \subseteq U$. We say a collection of multicurves *fills* a subsurface U if U is the smallest isotopy class of subsurfaces that contains the collection of multicurves. Any collection of curves fills a unique subsurface which we call the *fill* of that collection. Note, since we are considering an annulus to contain its core curve, the fill of a single curve c is the annulus with core curve c. If a curve c is contained in U, and U is the annulus with core curve c, then $U - c = \varnothing$. Otherwise, $U - c$ is the isotopy class of the subsurface obtained by deleting the representative of c that is contained in a representative of U. We extend this to define $U - \mu$ for a multicurve that is contained in the subsurface U.

Two subsurfaces are *disjoint* if they have no connected components in common and have disjoint representative. If the subsurfaces U and V are not disjoint and neither contains the other, we say U and V *overlap*. A curve c and subsurface U are *disjoint* if the annulus with core curve c is disjoint from U. This extends to define the disjointness of a multicurve and a subsurface. If a multicurve μ is not disjoint from a subsurface U, then we say μ and U *intersect*. If the multicurve μ intersects a subsurface U, but is disjoint from ∂U, then some component of μ is contained in U and we let $\mu \cap U$ denote this subset of components of μ.

A *pants subsurface* is any subsurface homeomorphic to $S^n_{0,p}$ with $p + n = 3$. Given a subsurface U of S we call a multicurve that is contained in U and has maximal cardinality among multicurves on U a *pants decomposition* of U. Note, when U has any annular components, then every pants decomposition of U contains the core curve for each of the annular components of U.

The *mapping class group*, $\mathrm{MCG}(S)$, of a surface S is the group of diffeomorphism of S modulo isotopies that fix the marked points. If $z \in S$ is not a marked point, then $\mathrm{MCG}(S; z)$ is the subgroup of $\mathrm{MCG}(S^z)$ that fixes the marked point z. When $S \cong S^0_{g,p}$, then $\mathrm{MCG}(S; z) = \mathrm{MCG}(S^z)$. The *point pushing subgroup*, $\mathrm{Push}(S; z)$, is the subgroup of $\mathrm{MCG}(S; z)$ comprised of all pushes of the point z along loops in S based at z. The Birman exact sequence tells us that $\mathrm{Push}(S; z)$ is naturally isomorphic to $\pi_1(S)$ and is the kernel of the surjective map $\mathrm{MCG}(S; z) \to \mathrm{MCG}(S)$ induced by forgetting the marked point z.

2.5 Complexes of curves and subsurface projection

Let $S \cong S^n_{g,p}$. Given a non-annular subsurface $U \subseteq S$, the *curve complex*, $\mathcal{C}(U)$, of U is the flag simplicial complex whose vertices are curves on U and with an edge between two distinct curves if they are disjoint. If U is connected and has complexity 1, then all curves of U intersect, so we modify the definition of the edges to allow for an edge between any two curves that intersect minimally on U—twice if $(g, p+n) = (0, 4)$ and once if $(g, p+n) = (1, 1)$.

When U is an annulus, then $\mathcal{C}(U)$ has an alternative description and is often called the *annular complex* instead of the curve complex. To define $\mathcal{C}(U)$ when U is an annulus, fix a complete hyperbolic metric on S where the marked points are viewed as punctures. Let

$\widetilde{U}$ be the annular cover of S corresponding to the annulus U equipped with the lift of this hyperbolic metric on S. Let $\widehat{U}$ denote the compactification of $\widetilde{U}$ to a closed annulus obtained from the usual compactification of the hyperbolic plane. The vertex set of $\mathcal{C}(U)$ is the set of isotopy classes (relative the boundary at infinity of $\widehat{U}$) of arcs that connect the two boundary components of $\widehat{U}$. Two vertices of the annular complex $\mathcal{C}(U)$ are then joined by an edge if they have representatives with disjoint interiors.

Masur and Minsky proved that $\mathcal{C}(U)$ is hyperbolic in all cases where it is non-empty [MM99]. A number of other authors have since proved that the constant of hyperbolicity is independent of the surface S or subsurface U [Aou13, Bow14, CRS14, HPW15].

For each subsurface $U \subseteq S$, Masur and Minsky defined a subsurface projection map $\pi_U \colon \mathcal{C}^0(S) \to 2^{\mathcal{C}(U)}$. We direct the reader to [MM00] for the full definition of the projection maps and instead recall the properties we will need for the present work.

Lemma 2.4 ([MM00, Lemma 2.2, 2.3])**.** *Let c and c' be curves on a surface S and U be a non-pants subsurface of S.*

- *$\pi_U(c) \neq \varnothing$ if and only if c intersects U and $c \neq \partial U$ when U is an annulus.*
- *$\operatorname{diam}(\pi_U(c)) \leqslant 3$ whenever $\pi_U(c) \neq \varnothing$.*
- *If c and c' are disjoint, then $d_{\mathcal{C}(U)}(\pi_U(\alpha), \pi_U(\beta)) \leqslant 2$ whenever $\pi_U(c)$ and $\pi_U(c')$ are both non-empty. If additionally $c' \subseteq U$, then every curve in $\pi_U(c)$ will not intersect $c' = \pi_U(c')$.*

Notation 2.5. When c and c' are two curves with $\pi_U(c) \neq \varnothing$ and $\pi_U(c') \neq \varnothing$, we use $d_U(c, c')$ to denote $d_{\mathcal{C}(U)}(\pi_U(c), \pi_U(c'))$. As before this is the minimum distance between the sets $\pi_U(c)$ and $\pi_U(c')$.

Distances in the curve complex are known to be bounded above in terms of the intersection number of curves; for example [Bow06]. This means we can modify the edge relation in $\mathcal{C}(U)$ to allow for a uniformly bounded number of intersections between curves without changing the quasi-isometry type of $\mathcal{C}(U)$. In our case, it will sometimes be convenient to work with the following variant of the curve complex that is quasi-isometric to $\mathcal{C}(U)$ and hence uniformly hyperbolic.

Definition 2.6 (Modified curve graph)**.** For a non-annular, non-pants subsurface $U \subseteq S$, define $\mathcal{C}'(U)$ to be the graph with has vertex set $\mathcal{C}^0(U)$ and with an edge between curves if they are either disjoint or intersect at most 4 times inside of a connected complexity 1 subsurface of S.

When S is a closed surface and $z \in S$, the map $\Pi \colon S^z \to S$ induced by forgetting the marked point z has the property that the image of every essential curve on S^z is still essential on S. Thus, Π induces a map $\Pi \colon \mathcal{C}^0(S^z) \to \mathcal{C}^0(S)$. A theorem of Kent, Leininger, and Schleimer proves that the preimage of a curve under Π is a tree.

Theorem 2.7 ([KLS09, Theorem 7.1])**.** *Let $S \cong S^0_{g,0}$ with $g \geqslant 2$. For each $\alpha \in \mathcal{C}^0(S)$, the subset of $\mathcal{C}(S^z)$ spanned by $\Pi^{-1}(\alpha)$ is a tree.*

The *pants graph*, $\mathcal{P}(S)$, of a surface S is the graph whose vertices are all pants decomposition of S and where two pants decompositions ρ and ρ' are joined by an edge if there exists curves $c \in \rho$ and $c' \in \rho'$ such that $(\rho - c) \cup c' = \rho'$ and c and c' are joined by an edge in the curve complex of the complexity 1 component of $S - (\rho - c)$. Pictorially, edges in $\mathcal{P}(S)$ correspond to one of the two "flip" moves shown in Figure 1.

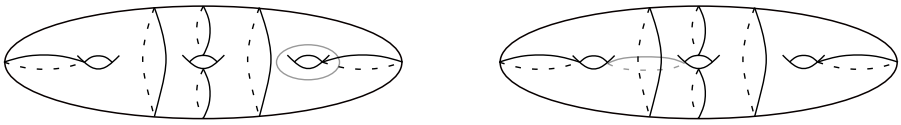

Figure 1: Replacing the gray curve in each picture with the black curve it intersects gives an example of the two different types of flip moves in the pants graph.

A *marking*[†] μ of a surface S is a set $\{(c_1, t_1), \ldots, (c_k, t_k)\}$ where $c_1, \ldots, c_k$ are curves that make a pants decomposition of S and each t_i is a vertex of the annular complex for the annulus whose core curve is c_i. We use $\text{base}(\mu)$ to denote the set $\{c_1, \ldots, c_k\}$ and $\text{markarc}(\mu)$ to denote $\{t_1, \ldots, t_k\}$. A marking μ is *clean* if for each $i \in \{1, \ldots, k\}$, the arc t_i is the lift of a curve c_i' such that c_i' and c_i fill a complexity 1 subsurface U and $d_U(c_i', c_i) = 1$. Given a marking μ and a non-pants subsurface U, Masur and Minsky defined the subsurface projection of μ as follows:

- If U is an annulus with core curve $c_i \in \text{base}(\mu)$, then $\pi_U(\mu) = t_i$.
- If U is any other subsurface, then $\pi_U(\mu) = \pi_U(\text{base}(\mu))$.

Masur and Minsky defined the *marking graph* $\mathcal{M}(S)$. This graph has the set of all clean markings of S as vertices, with edges defined by a "twist move" and a "flip plus clean-up" move. The next result summarizes the properties of the marking graph that we will need.

Theorem 2.8 ([MM00])**.** *The marking graph $\mathcal{M}(S)$ is a proper connected graph with the following properties:*

1. *The action of* $\text{MCG}(S)$ *on S induces a metrically proper and cocompact action of* $\text{MCG}(S)$ *on $\mathcal{M}(S)$ by isometries.*
2. *There is a function $f\colon [0, \infty) \to [0, \infty)$ depending only on S so that for any two clean markings $\mu, \nu \in \mathcal{M}(S)$ if $d_U(\mu, \nu) \leqslant r$ for each non-pants subsurface $U \subseteq S$, then $d_{\mathcal{M}(S)}(\mu, \nu) \leqslant f(r)$.*
3. *There is $D > 0$ depending only on S so that for each (not necessarily clean) marking μ, there exists a set of clean markings* $\text{cl}(\mu)$ *satisfying:*
   - $\text{base}(\mu') = \text{base}(\mu)$ *for each* $\mu' \in \text{cl}(\mu)$*;*
   - *if* $\mu' \in \text{cl}(\mu)$*, then for each* $(c_i, t_i) \in \mu$ *and* $(c_i, t_i') \in \mu'$*, we have* $d_{U_i}(t_i, t_i') \leqslant 3$ *where U_i is the annulus with core curve c_i;*

[†]What we call a marking, Masur and Minsky called a complete marking. As we will have no need for incomplete markings, we forgo the distinction.

- $\mathrm{cl}(\mu)$ *has diameter at most* D *in* $\mathcal{M}(S)$.

The set of clear markings $\mathrm{cl}(\mu)$ defined in Item 3 of Theorem 2.8 are called the *clean markings compatible with* μ. Item 2 is a special case of Masur and Minsky's celebrated distance formula for the marking graph; see [MM00, Theorem 6.12].

2.6 Combinatorial HHSs

Behrstock, Hagen, and Sisto's original definition of a hierarchically hyperbolic space required the construction of a large number of hyperbolic spaces and the verification of nine axioms [BHS19]. Recently, Behrstock, Hagen, Martin, and Sisto introduced *combinatorial hierarchically hyperbolic spaces* which reduces the construction of a hierarchically hyperbolic space to the construction of a pair of spaces (X, W) [BHMS20]. As we will not need to work directly with the full definition of a hierarchically hyperbolic space, we will forgo giving the definition and instead describe the combinatorial HHS machinery.

The starting place for a combinatorial hierarchically hyperbolic space is a pair of spaces (X, W) where X is a flag simplicial complex and W is an graph whose vertices are all the maximal simplices of X. We call such a graph W an *X-graph*. To help illuminate the definitions we will maintain a running example where X will be the curve complex $\mathcal{C}(S)$ and W is the pants graph $\mathcal{P}(S)$; $\mathcal{P}(S)$ is a $\mathcal{C}(S)$-graph as pants decompositions of S are exactly the vertices of the maximal simplices in $\mathcal{C}(S)$.

The definition of a combinatorial HHS (X, W) includes a number of properties that the pair (X, W) need to satisfy. Stating these properties requires a bit of set up. First, we need the following augmented version of X.

Definition 2.9 (Augmented graph)**.** If W is an X-graph for the flag simplicial complex X, we define the *W-augmented graph* X^{+W} as the graph with the same vertex set as X and with two types of edges:

1. (X-edge) If two vertices $x_1, x_2 \in X^0$ are joined by an edge in X, then x_1 and x_2 are joined by an edge in X^{+W}.

2. (W-edge) If Δ_1 and Δ_2 are maximal simplices of X that are joined by an edge in W, then each vertex of Δ_1 is joined by an edge to each vertex of Δ_2 in X^{+W}.

Example 2.10. For the case of $(\mathcal{C}(S), \mathcal{P}(S))$, the augmented graph $\mathcal{C}(S)^{+\mathcal{P}(S)}$ is a copy of the 1-skeleton of $\mathcal{C}(S)$ with additional edges between curves that intersect minimally in a complexity 1 subsurface.

Next we define an equivalence relation among simplices of X.

Definition 2.11. Let Δ and Δ' be simplices of the flag simplicial complex X. We write $\Delta \sim \Delta'$ if $\mathrm{lk}(\Delta) = \mathrm{lk}(\Delta')$. We define the *saturation* of Δ, $\mathrm{Sat}(\Delta)$, to be the set of vertices of X contained in a simplex in the $\sim$-equivalence class of Δ. That is $x \in \mathrm{Sat}(\Delta)$ if and only if there exists $\Delta' \sim \Delta$ so that x is a vertex of Δ'.

Example 2.12. If Δ is a simplex of $\mathcal{C}(S)$, then Δ defines a subsurface U_Δ that is the disjoint union of all the non-pants components of $S - \Delta$. The link of Δ in $\mathcal{C}(S)$ is then spanned by all curves on U_Δ, that is, $\mathrm{lk}(\Delta) = \mathcal{C}(U_\Delta)$. Two simplices Δ and Δ' will therefore have equal links if and only if $U_\Delta = U_{\Delta'}$. Further, the saturation of Δ will be the set of curves of on $S - U_\Delta$ plus the curves in ∂U_Δ, that is, $\mathrm{Sat}(\Delta) = \mathcal{C}^0(S - U_\Delta) \cup \partial U_\Delta$. This is because every non-peripheral curve on $S - U_\Delta$ is part of a pants decomposition of $S - U_\Delta$, and the join of this pants decomposition with ∂U_Δ produces a simplex Δ' with $U_{\Delta'} = U_\Delta$.

Finally we define two subspaces of X^{+W} that are associated to every simplex of X.

Definition 2.13. Let X be a flag simplicial complex and let W be an X-graph. For each simplex Δ of X, define

- Y_Δ to be $X^{+W} - \mathrm{Sat}(\Delta)$;
- $H(\Delta)$ to be the subgraph of Y_Δ spanned by the vertices in $\mathrm{lk}(\Delta)$.

For $H(\Delta)$, we are taking the link in X, not in X^{+W}, and then considering the subgraph of Y_Δ induced by those vertices. We give both Y_Δ and $H(\Delta)$ their intrinsic path metrics. By construction, we have $Y_\Delta = Y_{\Delta'}$ and $H(\Delta) = H(\Delta')$ whenever $\Delta \sim \Delta'$.

Example 2.14. As discussed in Example 2.12, for a simplex Δ of $\mathcal{C}(S)$, the saturation of Δ is the set of curves on $S - U_\Delta$ plus the multicurve ∂U_Δ. Thus $\mathcal{C}(S) - \mathrm{Sat}(\Delta)$ is spanned by the set of all curves that intersect U_Δ. The space Y_Δ is then $\mathcal{C}(S) - \mathrm{Sat}(\Delta)$ with extra edges between curves that intersect minimally in a complexity 1 subsurface of S. Because $\mathrm{lk}(\Delta) = \mathcal{C}(U_\Delta)$, the space $H(\Delta)$ is a copy of $\mathcal{C}(U_\Delta)$ with additional edges between pairs of curves that intersect minimally in a complexity 1 subsurface of U_Δ.

We can now state the definition of a combinatorial hierarchically hyperbolic space.

Definition 2.15 (Combinatorial HHS)**.** Let $\delta \geqslant 0$, X be a flag simplicial complex, and W be an X-graph. The pair (X, W) is a *δ-combinatorial HHS* if the following are satisfied.

1. If $\Delta_1, \ldots, \Delta_n$ are non-maximal simplices of X with $\mathrm{lk}(\Delta_1) \subsetneq \mathrm{lk}(\Delta_2) \subsetneq \cdots \subsetneq \mathrm{lk}(\Delta_n)$, then $n \leqslant \delta$.
2. The W-augmented graph X^{+W} is connected and δ-hyperbolic.
3. For each non-maximal simplex $\Delta \subseteq X$, the space $H(\Delta)$ is δ-hyperbolic and the inclusion $H(\Delta) \to Y_\Delta$ is a (δ, δ)-quasi-isometric embedding.
4. Whenever Δ_1 and Δ_2 are non-maximal simplices for which there exists a non-maximal simplex Λ such that $\mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\Delta_1) \cap \mathrm{lk}(\Delta_2)$ and $\mathrm{diam}(H(\Lambda)) \geqslant \delta$, then there exists a simplex Ω in the link of Δ_2 so that $\mathrm{lk}(\Delta_2 \bowtie \Omega) \subseteq \mathrm{lk}(\Delta_1)$ and all simplices Λ as above satisfy $\mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\Delta_2 \bowtie \Omega)$.
5. For each non-maximal simplex Δ and $x, y \in \mathrm{lk}(\Delta)$, if x and y are not joined by an X-edge of X^{+W} but are joined by a W-edge of X^{+W}, then there exits simplices Λ_x, Λ_y so that $\Delta \subseteq \Lambda_x, \Lambda_y$ and $x \bowtie \Lambda_x$, $y \bowtie \Lambda_y$ are vertices of W that are joined by an edge of W.

Remark 2.16. We will colloquially refer to conditions (1), (4), and (5) as the combinatorial conditions and conditions (2) and (3) as the geometric conditions.

Example 2.17. For $(\mathcal{C}(S), \mathcal{P}(S))$, the spaces $H(\Delta)$ are hyperbolic because they are quasi-isometric to $\mathcal{C}(U_\Delta)$. This is because distances in $\mathcal{C}(U_\Delta)$ are bounded above by the intersection number of curves. Because the space Y_Δ is spanned by curves that intersect the subsurface U_Δ, subsurface projection gives a coarsely Lipschitz map $\pi_{U_\Delta} \colon Y_\Delta \to H(\Delta)$. This makes $H(\Delta)$ quasi-isometrically embedded in Y_Δ by Lemma 2.3.

For the combinatorial conditions, observe that $\mathrm{lk}(\Delta) \subseteq \mathrm{lk}(\Delta')$ if and only if $U_\Delta \subseteq U_{\Delta'}$. Thus, the finite complexity of S gives a uniform bound on chains of properly nested links of simplices. For Item 4, it suffices to verify the condition for simplices Δ_1 and Δ_2 where U_{Δ_1} and U_{Δ_2} overlap. Let Q be the largest subsurface contained in both U_{Δ_1} and U_{Δ_2}, then let Γ be a simplex so that $\Delta_2 \bowtie \Gamma$ is a pants decomposition of $S - Q$, The desired simplex Ω will be Γ joined with the set of boundary curves of Q that are not curves of Δ_2. For Item 5, if x and y are two curves joined by a $\mathcal{P}(S)$-edge but not a $\mathcal{C}(S)$-edge, then x and y fill a complexity 1 subsurface V. If x and y are in the link of Δ, then $V \subseteq U_\Delta$. Hence, we can find Λ_x and Λ_y by taking the join of Δ with ∂U_Δ and pants decompositions of U_Δ that contain x and y respectively.

When (X, W) is a combinatorial HHS, the space W is a hierarchically hyperbolic space. Further, a group G will be hierarchically hyperbolic if it acts as described below on both X and W. In Section 7, we provide a brief summary of the salient parts of the definition of hierarchical hyperbolicity as well as a description of the hierarchically hyperbolic structure imparted on combinatorial HHSs.

Theorem 2.18 ([BHMS20, Theorem 1.18, Remark 1.19])**.** *If (X, W) is a δ-combinatorial HHS, then W is connected and a hierarchically hyperbolic space. Further, a finitely generated group G will be a hierarchically hyperbolic group if*

1. *G acts on X by simplicial automorphisms with finitely many orbits of links of simplices;*
2. *the action of G on maximal simplices of X induces a metrically proper and cobounded action of G on W by isometries.*

Remark 2.19. We emphasis that there is no a priori requirement that either X or W is connected. In the proof of Theorem 2.18, it is shown that the definition of a combinatorial HHS implies W is connected; see Sections 1 and 5.2 of [BHMS20]. We will take advantage of this in Section 5 to skip a direct proof of connectedness for our combinatorial HHS.

3 A sketch of the proof

Before embarking on the work for our main result, we give an outline of the proof in the case where α is a single curve on S. This case captures all of the key ideas of our proof while avoiding some technicalities that arise from α having multiple components. Recall, S^z is the surface obtained from S by adding z as a marked point of S. Let G_α be the stabilizer of α in $\mathrm{MCG}(S)$ and E_α be the preimage of G_α in $\mathrm{MCG}(S; z)$. Recall, $\Pi \colon S^z \to S$ is the map that forgets the marked point, and it induces a map $\Pi \colon \mathcal{C}^0(S^z) \to \mathcal{C}^0(S)$ because S is closed.

Our first step is to build a combinatorial HHS (X_α, W_α) where the vertex stabilizers of E_α on W_α are each generated by Dehn twists about disjoint curves. Let F_α be the subset of $\mathcal{C}(S^z)$ spanned by $\Pi^{-1}(\alpha)$. Theorem 2.7 established that F_α is a tree. While F_α is a hyperbolic space that E_α acts on, this action only "sees" the action of the $\pi_1(S) \cong \mathrm{Push}(S; z)$ subgroup of E_α; we need our hyperbolic space X_α to "see" the action of all of E_α. To accomplish this, we define X_α to be the subset of $\mathcal{C}(S^z)$ spanned by the preimages of all curves in the star of α in $\mathcal{C}(S)$. That is $X_\alpha^0 = \{c \in \mathcal{C}^0(S^z) : \Pi(c) \text{ is disjoint from or equal to } \alpha\}$.

The rationale for this choice of X_α is that the star of α in $\mathcal{C}(S)$ plus the pants graph on the surface $S - \alpha$ make a combinatorial HHS that sees all of the action of G_α except for Dehn twists. Thus, by including the preimage of the whole star of α, we capture both the $\pi_1(S)$ portion and the G_α portion of E_α.

The key lemma we prove about X_α is that it is F_α-*guided*. This means for each vertex $x \in X_\alpha^0 - F_\alpha$, the set $\mathrm{lk}(x) \cap F_\alpha$ is a subtree T_x of F_α and if x, y are adjacent vertices of X_α, then $T_x \cap T_y$ is non-empty. This allows for points in X_α to be connected by well behaved paths in F_α instead of X_α, and allows us to prove that X_α is a quasi-tree using a variant of Manning's bottleneck criteria.

The maximal simplices of X_α are pants decomposition of S^z that use only curves in X_α. Our X_α-graph W_α is therefore defined analogously to the pants graph: the vertices are pants decomposition using only curves in X_α with an edge between two pants decompositions that differ by flipping one curve to another curve that intersects it at most 4 times. The choice of intersection number 4 to define the edges is because 4 is the smallest number of intersections that two distinct element of F_α can have inside of a 4-punctured sphere subsurface.

For a multicurve $\mu \subseteq X_\alpha$, we define a subsurface U_μ to be the subsurface filled by the curves in the X_α-link of μ. Unlike in the case of the pants graph, U_μ can have annular components. Because the edges of X_α correspond to disjointness of curves, the nesting of links in X_α can then be translated into the nesting of these U_μ subsurfaces in S^z. Thus, the combinatorial conditions for (X_α, W_α) to be a combinatorial HHS follow from the topology of the U_μ subsurfaces similarly to how they are proved in the case of the pants graph.

For the geometric conditions, we split the links of multicurves into two cases. If μ is a simplex for X_α so that $\alpha \subseteq \Pi(\mu)$, then every curve on the subsurface $S^z - \mu$ is a curve in X_α. Thus, the link of μ in X_α will be the entire curve complex $\mathcal{C}(S^z - \mu)$. This means we can prove the quasi-isometric embedding and hyperbolicity conditions using the subsurface projection map $\pi_{S^z - \mu}$. On the other hand, if $\alpha \nsubseteq \Pi(\mu)$, then not every curve on $S^z - \mu$ will be a curve in X_α and the link of μ in X_α is some proper subcomplex of $\mathcal{C}(S^z - \mu)$. This makes using the subsurface projection map $\pi_{S^z - \mu}$ problematic, as we might get a curve that is not in X_α when we project a curve in $X_\alpha - \mathrm{Sat}(\mu)$ to $S^z - \mu$. Instead, we show that $X_\alpha - \mathrm{Sat}(\mu)$ is still guided by the fiber F_α because $\alpha \nsubseteq \Pi(\mu)$. This implies $\mathrm{lk}(\mu)$ will be a quasi-tree like X_α, and allows us to build a coarsely Lipschitz map onto $\mathrm{lk}(\mu)$ by first mapping a curve $x \in X_\alpha^0 - \mathrm{Sat}(\mu)$ into the fiber F_α and then taking a closest point projection in F_α to $\mathrm{lk}(\mu) \cap F_\alpha$.

Once we have proved (X_α, W_α) is a combinatorial HHS, we need to modify both spaces to produce a combinatorial HHS where E_α has a metrically proper action. Since the E_α-stabilizer of a vertex of W_α is generated by the Dehn twists around the curves in the pants decomposition, we need to modify X_α to account for these Dehn twists. We do this by "blowing up" each vertex of X_α as follows: for each curve $c \in X_\alpha$, we add as additional

vertices the set $B(c)$ of vertices of the annular complex for the annulus with core curve c. We add an edge between c and each element of $B(c)$ as well as between each element of $B(c)$ and every element of $B(c') \cup \{c'\}$ whenever c and c' are disjoint. We use $B(X_\alpha)$ to denote this blow-up of X_α.

The maximal simplices of $B(X_\alpha)$ are now (not necessarily clean) markings on S^z, and we define the graph $B(W_\alpha)$ to have vertices all the markings that are maximal simplices of $B(X_\alpha)$ with edges corresponding to both twist and flip moves. Since the vertices of $B(W_\alpha)$ are markings instead of pants decompositions, the action of E_α on $B(W_\alpha)$ will be metrically proper. At almost all steps, the proof that $(B(X_\alpha), B(W_\alpha))$ is a combinatorial HHS reduces to the proof that (X_α, W_α) is an combinatorial HHS. The only truly new conditions to be verified are the geometric conditions for the links of simplices μ where μ is a pants decomposition with all but one curve c marked. The hyperbolic space associated to such a link is quasi-isometric to the annular complex of the annulus with core curve c, thus we can use the subsurface projection onto this annulus to verify the quasi-isometric embedding condition.

4 Tree guided spaces

We now introduce the technique that we will use to verify that a graph or simplicial complex is a quasi-tree.

Definition 4.1. Let X be a graph and $F \subseteq X$ be a connected subgraph. For each $x \in X^0$ define $L(x) = \{x\}$ if $x \in F$ and $L(x) = \mathrm{lk}(x) \cap F$ if $x \notin F$. We say X is *F-guided* if

1. for each $x \in X^0$, $L(x)$ is a non-empty connected subset of F;
2. if $x, y \in X^0$ are joined by an edge, then either x and y are joined by an edge of F or $L(x) \cap L(y)$ is a non-empty connected subset of F.

When F is a tree and X is F-guided, then we say X is *tree guided*.

The key property imparted by Definition 4.1 is that every path in the graph X produces a connected subset of the guiding subgraph F connecting the end points of the path.

Definition 4.2. Let X be a graph and $F \subseteq X$ be a connected subgraph so that X is F-guided. If γ is an edge path in X with vertices $x_1, \ldots, x_n$, the sequence $L(x_1), \ldots, L(x_n)$ is called the *F-sequence* for γ. Definition 4.1 implies that for any edge path of X, the union of the $L(x_i)$ in the F-sequence span a connected subset of F.

The main result we will need is that tree guided graphs are quasi-isometric to trees.

Lemma 4.3. *Let X be a graph and $F \subseteq X$ be a connected subgraph. If X is F-guided and F is a tree, then X is uniformly quasi-isometric to a tree.*

To prove Lemma 4.3, we use the following variant of Manning's bottleneck criteria [Man05]. A proof of the this variant was given in Proposition 2.5 of [DDLS].

Proposition 4.4 ([Man05, DDLS])**.** *Let X be a graph. Suppose that there exists a constant R with the following property: for each pair of vertices $w, w' \in X$ there exists an edge path $\eta(w, w')$ from w to w' so that for any vertex v of $\eta(w, w')$, any path from w to w' in X intersects* $\mathrm{ball}_R(v)$*. Then X is quasi-isometric to a tree, with quasi-isometry constants depending on R only.*

Proof of Lemma 4.3. For each $x, y \in X^0$, let $\rho(x, y)$ be the shortest path in F connecting $L(x)$ to $L(y)$ in F. Define $\eta(x, y)$ to be the edge path spanned by $x \cup \rho(x, y) \cup y$. Let γ be any edge path of X connecting x and y. The F-sequence for γ gives a connected subset of F that contains the end points of $\rho(x, y)$. Since F is a tree, the F-sequence of γ must then contain all of $\rho(x, y)$. In particular, for each $v \in \rho(x, y)$, the path γ must intersect the ball of radius 1 around v. Hence, X is uniformly quasi-isometric to a tree by Proposition 4.4. □

We will also need the following lemma that will allow us to preserve F-guidedness when removing a set of vertices that are distinct from F.

Lemma 4.5. *Let X be a graph and $F \subseteq X$ be a connected subgraph so that X is F-guided. If $Z \subseteq X^0$ with $Z \cap F = \varnothing$, then $X - Z$ is F-guided*

Proof. Because $Z \cap F = \varnothing$, the set $L(x)$ is unchanged for each vertex $x \in X - Z$. Further, if x_1 and x_2 are adjacent vertices of $X - Z$, then they are also adjacent vertices of X and either $x_1, x_2 \in F$ or $L(x_1) \cap L(x_2)$ is a non-empty connected subset of F as desired. □

We conclude by recording a more general version of Lemma 4.3 that will we will not need, but may be useful for future work.

Lemma 4.6. *Let X be a graph and $F \subseteq X$ be a connected subgraph. If X is F-guided, then X is $(2, 0)$-quasi-isometric to the graph obtained from F by adding a vertex v_x for each $x \in X^0 - F^0$ and connecting v_x to every vertex of $L(x)$. Moreover, if F is hyperbolic and each $L(x)$ is uniformly quasiconvex in F, then X is hyperbolic.*

Proof. Let $\widehat{F}$ denote the graph obtained from F by adding a vertex v_x for each $x \in X^0 - F^0$ and connecting v_x to every vertex of $L(x)$. Let $f \colon \widehat{F} \to X$ be given by $f(w) = w$ for $w \in F$ and $f(v_x) = x$ for $x \in X^0 - F^0$.

We first prove that f is a quasi-isometry. Let $z_1 = w_0, w_1, \ldots, w_n = z_2$ be the vertices of the $\widehat{F}$-geodesic connecting $z_1, z_2 \in \widehat{F}^0$. Since X is F-guided, $f(w_0), f(w_1), \ldots, f(w_n)$ are the vertices of a path in X connecting $f(z_1)$ and $f(z_2)$. Thus, $d_X(f(z_1), f(z_2)) \leqslant d_{\widehat{F}}(z_1, z_2)$ for all $z_1, z_2 \in \widehat{F}$. For the other inequality, assume $f(z_1)$ and $f(z_2)$ are joined by an edge of X while z_1, z_2 are not joined by an edge of $\widehat{F}$. This only occurs if $z_1 = v_x$ and $z_2 = v_y$ for $x, y \in X^0 - F^0$. Since $f(z_1) = v_x$ and $f(z_2) = v_y$ are joined by an edge of X, the second condition in the definition of a F-guided space ensures that $L(x) \cap L(y) \neq \varnothing$. But this means $d_{\widehat{F}}(x, y) = 2$. Hence, we have $d_{\widehat{F}}(z_1, z_2) \leqslant 2d_X(f(z_1), f(z_2))$ for all $z_1, z_2 \in \widehat{F}$.

When F is hyperbolic and each $L(x)$ is uniformly quasiconvex in F, Proposition 2.6 of [KR14] shows that $\widehat{F}$ is hyperbolic with constant depending only on the hyperbolicity of F and the quasiconvexity constant of the $L(x)$'s. Since X is uniformly quasi-isometric to $\widehat{F}$, X must be hyperbolic as well. □

5 A combinatorial HHS without annuli

Fix a closed surface $S \cong S^0_{g,0}$ with $g \geqslant 2$ and let S^z be the surface obtained from S by making $z \in S$ a marked point. Fix a multicurve α on S and let G_α be the stabilizer of α in $\mathrm{MCG}(S)$. Let E_α be the full preimage in $\mathrm{MCG}(S; z)$ of G_α. Recall that the map $\Pi\colon \mathcal{C}^0(S^z) \to \mathcal{C}^0(S)$ is the map on curves induced by removing the marked point of S^z.

In this section, we build a combinatorial HHS (X_α, W_α) on which E_α acts with large vertex stabilizers. In Section 6, we will modify (X_α, W_α) to produce a combinatorial HHS denoted $(B(X_\alpha), B(W_\alpha))$ where E_α does have a metrically proper action on $B(W_\alpha)$. As we can reduce the proof for $(B(X_\alpha), B(W_\alpha))$ to the case of (X_α, W_α), this intermediate step allows us to present a simpler and more transparent proof. Further, the spaces X_α and W_α are of intrinsic interest as they are the most direct analogues of the curve complex and pants graph for E_α.

Hence forth, we say that constants are *uniform* if they do not depend on the surface S, the multicurve α, or a simplex of X_α. A *uniform quasi-tree* is thus a quasi-tree where the quasi-isometry constants to a tree does not depend on any of these quantities.

5.1 Definition of X_α and W_α

We define the hyperbolic space X_α by taking the preimage of the star of the multicurve α under the marked point forgetting map $\Pi\colon \mathcal{C}^0(S^z) \to \mathcal{C}^0(S)$.

Definition 5.1. Define X_α to be the subgraph of $\mathcal{C}(S^z)$ spanned by the curves c so that $\Pi(c)$ is either disjoint from or contained in α on S. If $\alpha_1, \ldots, \alpha_m$ are the curves comprising α, then define F_i to be the subset of $X_\alpha \subseteq \mathcal{C}(S^z)$ spanned by $\Pi^{-1}(\alpha_i)$. We call F_i the *fiber over* α_i. By Theorem 2.7, F_i is a tree for each $i \in \{1, \ldots, m\}$.

The simplices of X_α are precisely the multicurves on S^z whose components are all curves in X_α. Thus, we will freely go between a multicurve of curves in X_α and the simplex its vertices span.

We will exploit the fact that the fibers F_i are trees to prove X_α is itself a quasi-tree. Our key tool is the fact that whenever a simplex μ does not intersect a fiber F_i, then the link of μ intersects F_i in a subtree of F_i.

Lemma 5.2. *Let μ be a simplex in X_α, and suppose there exists a fiber F_i so that $\mu \cap F_i = \varnothing$. For all $a_1, a_2 \in \mathrm{lk}(\mu) \cap F_i$, there is a path in $\mathrm{lk}(\mu) \cap F_i$ connecting a_1 to a_2.*

Proof. Let $\mu \subseteq X_\alpha$ be a simplex and F_i be a fiber so that $\mu \cap F_i = \varnothing$. This means $\Pi(\mu)$ is disjoint from α_i on S ensuring that $\mathrm{lk}(\mu) \cap F_i \neq \varnothing$. Let $a_1, a_2 \in \mathrm{lk}(\mu) \cap F_i$. We prove by induction on the intersection number $i(a_1, a_2)$ that a_1 is connected to a_2 in $\mathrm{lk}(\mu) \cap F_i$.

The base case of $i(a_1, a_2) = 0$ is automatic since $i(a_1, a_2) = 0$ implies a_1 is joined to a_2 by an edge of F_i and this edge is in $\mathrm{lk}(\mu) \cap F_i$ since $a_1, a_2 \in \mathrm{lk}(\mu)$.

Assume that the lemma holds for curves of $\mathrm{lk}(\mu) \cap F_i$ that intersect strictly less than n times and suppose $0 < i(a_1, a_2) = n$. Since a_1 and a_2 are not disjoint and are both elements of F_i, we know $a_1 \cup a_2$ must form a bigon that contains the marked point of S^z. In particular, there exists an innermost bigon of $a_1 \cup a_2$ around the marked point. We can surger a_2 across

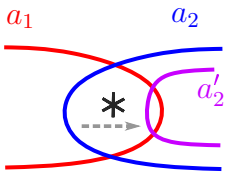

Figure 2: The surgery of a curve across a bigon containing the marked point.

this inner most bigon as shown in Figure 2 to produce a curve a_2' so that a_2' is disjoint from a_2 and $i(a_1, a_2') < n$. Now, a_2' is disjoint from μ on S^z, since a_2 is disjoint from μ on S^z. Further, $a_2' \in F_i$ since a_2' would be isotopic to a_2 if the marked point of S^z was removed. Thus, $a_2' \in \mathrm{lk}(\mu) \cap F_i$ with $i(a_1, a_2') < n$ and the induction hypothesis implies the path in F_i from a_1 to a_2' is contained in $\mathrm{lk}(\mu)$. Since $i(a_2, a_2') = 0$, this gives a path in $\mathrm{lk}(\mu) \cap F_i$ from a_1 to a_2. □

Using Lemma 5.2, we can prove X_α is a quasi-tree by proving it is F_i-guided for each fiber F_i.

Proposition 5.3. *For each $i \in \{1, \dots, m\}$, the graph X_α is F_i-guided. In particular, X_α is a uniform quasi-tree.*

Proof. Since each fiber F_i is a tree, the final clause follows from the first clause by Lemma 4.3.

Fix a fiber F_i over α_i. For $x \in X_\alpha^0$, define $L_i(x) = \{x\}$ if $x \in F_i$ and $L_i(x) = \mathrm{lk}(x) \cap F_i$ if $x \notin F_i$. By definition of X_α, $L_i(x)$ is a non-empty subset of F_i for all $x \in X_\alpha^0$, and by Lemma 5.2, $L_i(x)$ is connected. We now verify the second requirement of being F_i-guided: that x being adjacent to y in X_α implies $x, y \in F_i$ or $L_i(x) \cap L_i(y)$ is non-empty and connected.

Let $x, y \in X_\alpha^0$ be joined by an edge in X_α. Thus, x and y are disjoint curves on S^z. Without loss of generality, assume $x \notin F_i$. If $y \in F_i$, then $y \in L_i(x)$ by definition and we are done because $L_i(x) \cap L_i(y) = \{y\}$. We can therefore assume $y \notin F_i$ as well. Similar to the proof of Lemma 5.2, we will use induction on intersection number to prove that there exists a curve a in $L_i(x) \cap L_i(y)$. Since F_i is a tree and $L_i(x)$ and $L_i(y)$ are already connected, this suffices to prove $L_i(x) \cap L_i(y)$ is connected.

Let a be any curve in $L_i(x)$. If $i(a, y) = 0$, then $a \in L_i(y)$ by definition, so suppose $i(a, y) > 0$. Since $a \in F_i$ and $y \notin F_i$, we must have that $\Pi(y)$ and $\Pi(a)$ are disjoint on S. Thus, y and a must form an innermost bigon that contains the marked point of S^z. We can surger a across the marked point, to produce a curve a' with $i(a', y) < i(a, y)$ and a' disjoint from x (this surgery is identical to the surgery in Figure 2 with $y = a_1$, $a_2 = a$ an $a_2' = a'$). Since a' is obtained from a by surgery across the marked point, $\Pi(a) = \Pi(a')$. This means $a' \in F_i$ and we have $a' \in L_i(x)$. Since $i(a', y) < i(a, y)$, we can now induct on the intersection number to find $a'' \in L_i(x)$ with $i(a'', y) = 0$. Such a curve a'' will be in $L_i(y)$ by the definition of edges in X_α. □

The next lemma verifies that the maximal simplices of X_α are exactly the pants decomposition of S^z consisting of curves in X_α^0. We give a more technical version of this fact that we will need later.

Lemma 5.4. *Let Q be any subsurface of S^z and γ be a (possibly empty) multicurve on Q. If both γ and ∂Q are simplices of X_α, then there exists a pants decomposition τ of Q so that $\gamma \subseteq \tau \subseteq X^0_\alpha$.*

Proof. First we show that we can assume Q has no annular components without losing any generality. Since $\partial Q \subseteq X_\alpha$, the core curve of every annular component of Q is a curve of X^0_α. Thus, if Q_0 is the union of the non-annular components of Q and $\tau_0 \subseteq X_\alpha$ is a pants decomposition of Q_0 containing $\gamma \cap Q_0$, then we can define τ to be the union of τ_0 and the core curves of the annular components of Q. Hence, we can assume Q has no annular component.

In addition to assuming Q has no annular component, we can also assume Q has no pants components as pairs of pants do not contain any curves. These two assumptions mean that $\Pi(Q)$ does not contain any annular components.

If $\Pi(Q) \subseteq S - \alpha$, then every curve on Q must be a curve in X_α. This means any pants decomposition of Q that contains γ is a simplex of X_α.

Assume then that $\Pi(Q) \nsubseteq S - \alpha$. Since ∂Q is a simplex of X_α, no component of α will intersect $\partial\Pi(Q) = \Pi(\partial Q)$. Because $\Pi(Q)$ has no annular components, the only way for $\Pi(Q) \nsubseteq S - \alpha$ without α intersecting $\partial\Pi(Q)$ is for a curve of α to be contained in $\Pi(Q)$.

If Q does not contain the marked point z, then Π restricted to Q is a homeomorphism. Hence, for each curve $\alpha_i \in \alpha \cap \Pi(Q)$ there is exactly one curve a_i on Q with $\Pi(a_i) = \alpha_i$. Thus, there is a multicurve η that contains all of the curves of $\Pi^{-1}(\alpha)$ that are on Q. Moreover, every curve of $Q - \eta$ is a curve in X_α. Therefore, any pants decomposition of Q that contains η will be a simplex of X_α. Since $\gamma \subseteq X_\alpha$, γ cannot intersect η. Hence any pants decomposition τ that contains $\gamma \cup \eta$ will suffice.

Finally, assume Q does contain the marked point z. If $Q - \gamma$ has a pants component P that contains the marked point, then Π restricted to $Q - P$ is a homeomorphism. Hence, the same reasoning as the previous paragraph yields a pants decomposition τ' of $Q - P$ that is a simplex of X_α and contains $\gamma \cap (Q - P)$. The desired simplex is then $\tau = \tau' \cup \gamma$. If $Q - \gamma$ does not have a pants component that contains the marked point, then for some $\alpha_i \in \alpha \cap \Pi(Q)$ there exist $a_1, a_2 \in \Pi^{-1}(\alpha_i)$ so that $a_1, a_2 \subseteq Q$ and $Q - (a_1 \cup a_2)$ has a pants component P that contains the marked point. We can additionally choose a_1 and a_2 so that every curve of γ is either equal to or disjoint from each a_i. As before, there is a pants decomposition τ' of $Q - P$ that is a simplex of X_α and contains $\gamma \cap (Q - P)$. In this case $\tau = \tau' \cup a_1 \cup a_2$ is the desired pants decomposition. □

We now define the X_α-space W_α.

Definition 5.5. Define W_α to be the the following graph.

- Vertices: Maximal simplices of X_α.
- Edges: Distinct vertices ρ and ρ' are joined by an edge if there exists curves $c \in \rho$ and $d \in \rho'$ so that $i(c, d) \leqslant 4$ and $\rho = (\rho' - d) \cup c$.

By Lemma 5.4, the vertices of W_α are all pants decomposition of S^z that are built from curves in X_α. Two such pants decomposition are joined by an edge precisely when they differ by at most one curve that intersect at most 4 times. Hence, W_α is an analogue of the pants graph for the group E_α.

5.2 Combinatorial conditions for (X_α, W_α)

This subsection is devoted to verifying that (X_α, W_α) satisfies the combinatorial properties of the definition of combinatorial HHS (Items 1, 4, and 5 of Definition 2.15). We start by explaining how the combinatorics of links of simplices of X_α are encoded in the topology of subsurfaces of S^z. Through out this subsection, $\mathrm{lk}(\cdot)$ will denote links in the simplicial complex X_α.

Definition 5.6 (Subsurface for a multicurve)**.** Given a simplex μ of X_α, define U_μ to be the (possibly disconnected) subsurface filled by the curves of X_α that are disjoint from μ. That is, U_μ is the subsurface filled by the curves in $\mathrm{lk}(\mu)$.

The subsurface U_μ is automatically a subsurface of $S^z - \mu$, but U_μ might be strictly smaller than the collection of non-pants subsurfaces of $S^z - \mu$. To describe U_μ, let M_μ be the component of $S^z - \mu$ that contains the marked point of S^z then let α_μ be the set of curves $\{a \in \Pi^{-1}(\alpha) : a \subseteq (S^z - \mu) - M_\mu\}$. Since $(S^z - \mu) - M_\mu$ does not contain the marked point, α_μ is a multicurve on $S^z - \mu$. The next lemma asserts that U_μ is the union of the non-pants components of $S^z - (\mu \cup \alpha_\mu)$ and the annuli whose core curves are in α_μ. Figure 3 gives an example of U_μ that illustrates the lemma.

Lemma 5.7. *Let μ be a simplex of X_α. In the notation of the proceeding paragraph, U_μ is the disjoint union of the non-pants components of $S^z - (\mu \cup \alpha_\mu)$ plus the annuli whose core curves are in α_μ. Moreover, $\partial U_\mu \subseteq X_\alpha$ since every curve of ∂U_μ is either a curve of μ or a curve of α_μ.*

Proof. Let V_μ be the disjoint union of the non-pants components of $S^z - (\mu \cup \alpha_\mu)$ plus the annuli whose core curves are in α_μ. The curves in $\alpha_\mu \cup \mathrm{lk}(\mu \cup \alpha_\mu)$ fill V_μ. Since U_μ is filled by the curves in $\mathrm{lk}(\mu)$, it suffices to prove $\mathrm{lk}(\mu) = \alpha_\mu \cup \mathrm{lk}(\mu \cup \alpha_\mu)$. Since $\alpha_\mu \subseteq \mathrm{lk}(\mu)$ we have $\alpha_\mu \cup \mathrm{lk}(\mu \cup \alpha_\mu) \subseteq \mathrm{lk}(\mu)$. For the other direction, let $c \in \mathrm{lk}(\mu)$ and let K be the component of $S^z - \mu$ that contains c. If $K = M_\mu$, then c is disjoint from each curve in α_μ making $c \in \alpha_\mu \cup \mathrm{lk}(\mu \cup \alpha_\mu)$. If $K \neq M_\mu$, then Π restricted to K is a homeomorphism. Hence every curve of X_α that is on K must be either an element of $\alpha_\mu \cap K$ or disjoint from each curve in $\alpha_\mu \cap K$. Thus, either $c \in \alpha_\mu$ or $c \in \mathrm{lk}(\mu \cup \alpha_\mu)$. □

Using the subsurfaces U_μ, we can relate the combinatorics of the links in X_α to the topology of subsurfaces on S^z.

Lemma 5.8 (Topology to cHHS dictionary)**.** *Let μ and μ' be non-maximal simplices of X_α.*

1. *$\mathrm{lk}(\mu)$ is spanned by the set of all curves $\{x \in X_\alpha^0 : x \subseteq U_\mu\}$. In particular, $\mathrm{lk}(\mu)$ is a join whenever U_μ is disconnected.*
2. *$\mathrm{lk}(\mu) \subseteq \mathrm{lk}(\mu')$ if and only if $U_\mu \subseteq U_{\mu'}$.*
3. *$\mathrm{lk}(\mu) = \mathrm{lk}(\mu')$ if and only if $U_\mu = U_{\mu'}$.*
4. *$\mathrm{Sat}(\mu)$ is the set of all curves of X_α on $S^z - U_\mu$ plus the set of curves in ∂U_μ that are not core curves of any annular components of U_μ.*

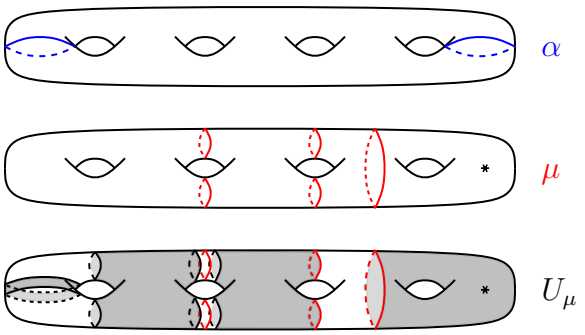

Figure 3: An example of the subsurface U_μ. The top picture shows the fixed multicurve α on the closed surface S. The second picture is of the multicurve $\mu \subseteq X_\alpha$ on the surface S^z (where $*$ denotes the marked point z). The bottom picture is the subsurface U_μ. The core curve of the annulus in the bottom picture is α_μ.

Proof. 1. The first item is a direct consequence of the fact that edges in X_α correspond to the disjointness of curves plus the definition of U_μ as the fill of $\mathrm{lk}(\mu)$.

2. Since U_μ and $U_{\mu'}$ are the subsurfaces filled by the curves in $\mathrm{lk}(\mu)$ and $\mathrm{lk}(\mu')$ respectively, $\mathrm{lk}(\mu) \subseteq \mathrm{lk}(\mu')$ implies $U_\mu \subseteq U_{\mu'}$.

   On the other hand, if $U_\mu \subseteq U_{\mu'}$ and c is a curve in $\mathrm{lk}(\mu)$, then $c \subseteq U_\mu \subseteq U_{\mu'}$. This makes $c \in \mathrm{lk}(\mu')$ by the first item.

3. This item is an immediate consequence of the second item.

4. Let τ be the subset of curves in ∂U_μ that are not core curves of any annular components of U_μ. By the first item, if a curve c is in $\mathrm{Sat}(\mu)$, then c must be disjoint from U_μ. Hence either $c \subseteq S^z - U_\mu$ or $c \in \tau$. To see that $\mathrm{Sat}(\mu)$ contains every curve $c \in X^0_\alpha$ that is a curve on $S^z - U_\mu$, observe that Lemma 5.4 provides a pants decomposition ρ of $S^z - U_\mu$ that contains c and uses only curves in X_α. Thus, $\rho \bowtie \tau$ is a simplex of X_α with $\mathrm{lk}(\rho \bowtie \tau) = \mathrm{lk}(\mu)$ because $U_{\rho \bowtie \tau} = U_\mu$. Hence $c \in \mathrm{Sat}(\mu)$. □

Armed with Lemma 5.8, we now verify the combinatorial criteria for (X_α, W_α) to be a combinatorial HHS.

Proposition 5.9. *Suppose μ and μ' are non-maximal simplices of X_α so that there exists a non-maximal simplex ν of X_α with $\mathrm{lk}(\nu) \subseteq \mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$ and $\mathrm{lk}(\nu)$ not a join or a single vertex. There exists a (possibly empty) simplex $\rho \subseteq \mathrm{lk}(\mu')$ so that $\mathrm{lk}(\mu' \bowtie \rho) \subseteq \mathrm{lk}(\mu)$ and if ν is any simplex as in the proceeding sentence, then $\mathrm{lk}(\nu) \subseteq \mathrm{lk}(\mu' \bowtie \rho)$.*

Proof. Assume μ and μ' are simplices of X_α as described in the proposition. Let V_0 be the subsurface of S^z filled by the curves in $\mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$, then let V be the disjoint union of the non-annular components of V_0. Note, if $V = \varnothing$, then $\mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$ is either empty or is a simplex of X_α. However, the assumption that there is a non-maximal simplex ν of X_α so that $\mathrm{lk}(\nu)$ is not a join or a single vertex and $\mathrm{lk}(\nu) \subseteq \mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$ prevents $\mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$ from being empty or a simplex. Thus, we know that $V \neq \varnothing$ and contains an infinite number of curves of X_α.

Suppose there is a simplex $\rho \subseteq \mathrm{lk}(\mu')$ with the property that $U_{\mu' \bowtie \rho} = V$ and let ν be a non-maximal simplex of X_α so that

1. $\mathrm{lk}(\nu) \subseteq \mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$;

2. $\mathrm{lk}(\nu)$ is not a join or a single vertex.

Then $U_\nu \subseteq V_0$ because $\mathrm{lk}(\nu) \subseteq \mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$. Since $\mathrm{lk}(\nu)$ is not a join, no curve in $\mathrm{lk}(\nu)$ can be a core curve for an annular component of V_0. Hence $U_\nu \subseteq V = U_{\mu' \bowtie \rho}$ which means $\mathrm{lk}(\nu) \subseteq \mathrm{lk}(\mu' \bowtie \rho)$ by Lemma 5.8. Since $U_{\mu' \bowtie \rho} = V \subseteq U_\mu$ implies $\mathrm{lk}(\mu' \bowtie \rho) \subseteq \mathrm{lk}(\mu)$, the simplex ρ would satisfy the proposition.

The proof will therefore be complete if we can find a simplex $\rho \subseteq \mathrm{lk}(\mu')$ so that $U_{\mu' \bowtie \rho} = V$. We will find ρ by applying Lemma 5.4 to the subsurface $U_{\mu'} - V$. In order to do so, we need to verify that the boundary curves of V are also curves in X_α.

Claim 5.10. ∂V is a simplex of X_α.

Proof. For the purposes of contradiction, assume $\partial V \nsubseteq X^0_\alpha$. This means there exists $a \in \Pi^{-1}(\alpha)$ so that $\Pi(a)$ and $\Pi(\partial V)$ intersect. Further, by performing surgery across the marked point as in Figure 2, we can pick a so that a and ∂V do not form any bigons around the marked point. Pick a representative of a that minimizes the intersection number with ∂V and let a' be the set of arcs produced by intersecting that representative with V.

First assume that V does not contain the marked point z. In this case, the forgetful map Π restricted to V is a homeomorphism. Thus no curve of $\mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$ on V can intersect a'. Hence $V - a'$ is a proper subsurface of V that contains all the curves of $\mathrm{lk}(\mu) \cap \mathrm{lk}(\mu')$ that are not the core curves of an annular component of V_0. However this is impossible as V is the subsurface filled by these curves. Hence, ∂V must be a simplex of X_α.

Now assume that V does contain the marked point z. Because a is not contained in V_0, it must be the case that a intersects either μ or μ'. Without loss of generality, assume a intersects μ. Since $\Pi(a)$ does not intersect $\Pi(\mu)$, a and μ must form an inner most bigon B around the marked point. Because V contains the marked point, the bigon B must intersect the subsurface V. Since a and ∂V do not intersect in any bigon around the marked point, the only way for B to intersect V is for μ to also intersects V. However, this is impossible because V is filled by curves from $\mathrm{lk}(\mu)$. Hence ∂V must be a simplex of X_α ■

Since both ∂V and $\partial U_{\mu'}$ are both simplices of X_α, the boundary of $U_{\mu'} - V$ is also a simplex of X_α. Hence, Lemma 5.4 provides a pants decomposition τ of $U_{\mu'} - V$ made up of curves of X_α. Define ρ to be the simplex spanned by $\tau \cup \partial V$.

Claim 5.11. $U_{\mu' \bowtie \rho} = V$.

Proof. Every curve on V is disjoint from $\mu' \cup \rho$ by construction. Since V is filled by the curves of X_α that it contains and $U_{\mu' \bowtie \rho}$ is the subsurface filled by $\mathrm{lk}(\mu' \bowtie \rho)$, this means $V \subseteq U_{\mu' \bowtie \rho}$.

For the other direction, consider $c \in \mathrm{lk}(\mu' \bowtie \rho)$. Such a curve must be contained in a component of $U_{\mu'}$ and must be disjoint from ρ. Since ρ contains both a pants decomposition of $U_{\mu'} - V$ and ∂V, the only way for that to happen is for $c \subseteq V$. This implies $U_{\mu' \bowtie \rho} \subseteq V$ since $U_{\mu' \bowtie \rho}$ is filled by the curves in $\mathrm{lk}(\mu' \bowtie \rho)$. ■

As described before Claim 5.10, Claim 5.11 completes the proof of Proposition 5.9. □

Proposition 5.12. *Let μ be a non-maximal simplex of X_α and let x, y be curves of X_α contained in $\mathrm{lk}(\mu)$. If x and y are joined by a W_α-edge but not an X_α-edge in $X_\alpha^{+W_\alpha}$, then there exists a multicurve ν so that $\mu \subseteq \nu$ and $\nu \cup x$ and $\nu \cup y$ are adjacent vertices of W_α.*

Proof. If $x, y \in X_\alpha^{+W_\alpha}$ are vertices that are joined be a W_α-edge but not an X_α-edge, then the curve x and the curve y fill a complexity 1 subsurface V of S^z. Further, $\partial V \subseteq X_\alpha$ because the curves of ∂V are part of the maximal simplices of X_α that contain x and y and are joined by an edge in W_α.

Since $x, y \in \mathrm{lk}(\mu)$, every curve in μ is either a boundary curve of V or contained in $S^z - V$. By Lemma 5.4, we can extend $\mu \cap (S^z - V)$ to a pants decomposition ρ of $S^z - V$ that contains only curves of X_α. Thus, $\nu = \rho \cup \partial V$ is the desired multicurve. □

5.3 Geometric conditions for (X_α, W_α)

We begin the verification of the geometric conditions of a combinatorial HHS by checking that $X_\alpha^{+W_\alpha}$ is uniformly quasi-isometric to X_α. By Proposition 5.3, this makes $X_\alpha^{W_\alpha}$ a quasi-tree.

Corollary 5.13. *$X_\alpha^{+W_\alpha}$ is uniformly E_α-equivariantly quasi-isometric to X_α, and hence is a uniform quasi-tree.*

Proof. Because $X_\alpha^{+W_\alpha}$ is a copy of X_α with additional edges, it suffices to verify that whenever two vertices x, y of $X_\alpha^{+W_\alpha}$ are joined by a W_α-edge, then the distance between x and y in X_α is uniformly bounded. If x and y are joined by a W_α-edge, then there exists a multicurve ν so that $\nu \cup x$ and $\nu \cup y$ are adjacent edges of W_α. As the vertices of W_α are maximal simplices in X_α, both x and y are joined by an X_α-edge to each curve of ν. Hence $d_{X_\alpha}(x, y) \leqslant 2$. □

For the remainder of this section, we analyze the links in X_α to verify the second geometric condition for (X_α, W_α) to be a combinatorial HHS (Item 3 of Definition 2.15). We begin by recalling the notation we set out in Section 2.6. Note, all links of simplices in this section are taken in the the complex X_α.

Recall, that if Z is a subset of vertices of a graph X, then $X - Z$ is the subgraph spanned by the vertices in X that are not in Z. For a simplex μ in X_α the space Y_μ is defined as the space $X_\alpha^{+W_\alpha} - \mathrm{Sat}(\mu)$ and $H(\mu)$ is the subset of Y_μ spanned by $\mathrm{lk}(\mu)$. We endow both Y_μ and $H(\mu)$ with their intrinsic path metrics and not the subspace metric from X_α. Over the next two subsections, we will prove that $H(\mu)$ is uniformly hyperbolic and that the inclusion of $H(\mu)$ into Y_μ is uniformly a quasi-isometric embedding.

The simplices of X_α come in two types and we split our arguments accordingly.

Definition 5.14. Let μ be a non-maximal simplex of X_α and $F_1, \ldots, F_m$ be the fibers over $\alpha_1, \ldots, \alpha_m$.

1. (Subsurface type) We say μ is of *subsurface type* if μ contains a vertex of each fiber F_i. In this case, $\mathrm{lk}(\mu)$ is precisely $\mathcal{C}(S^z - \mu)$, because any curve on S^z that is disjoint from μ will be disjoint from α in the image of Π. By Lemma 5.7, U_μ will have no annular components, ensuring that $\mathcal{C}(S^z - \mu) = \mathcal{C}(U_\mu)$ in this case.

2. (Tree guided type) We say μ is of *tree guided type* if there exists a fiber F_i so that μ does not contain a vertex of F_i. In this case, $\mathrm{lk}(\mu)$ is a proper subset of $\mathcal{C}(S^z - \mu)$.

5.3.1 Subsurface type links

We handle the case where μ is of subsurface type first. Because of the definition of W_α and the construction of $X_\alpha^{+W_\alpha}$, Proposition 5.12 implies that $H(\mu)$ is a copy of $\mathrm{lk}(\mu)$ with additional edges between curves that intersect at most four times in a connected, complexity 1 subsurface of S^z. Since $\mathrm{lk}(\mu) = \mathcal{C}(S^z - \mu)$, this makes $H(\mu)$ equal to $\mathcal{C}'(S^z - \mu)$, the modified curve graph from Definition 2.6. We can therefore use the subsurface projection map to prove $H(\mu)$ is quasi-metrically embedded in Y_μ.

Proposition 5.15. *If μ is a subsurface type simplex of X_α, then $H(\mu) = \mathcal{C}'(S^z - \mu)$. In particular, $H(\mu)$ is uniformly hyperbolic.*

Proof. The vertices of $H(\mu)$ are all of the curves of S^z that are contained in $S^z - \mu$. By Proposition 5.12, there is an edge between two vertices of $H(\mu)$ if the curves are either disjoint or if they intersect at most four times inside of a complexity 1 subsurface of $S^z - \mu$. However, this is precisely the definition of $\mathcal{C}'(S^z - \mu)$. □

Proposition 5.16. *If μ is a subsurface type simplex of X_α, then the inclusion of $H(\mu)$ into Y_μ is a uniform quasi-isometric embedding.*

Proof. Because $Y_\mu = X_\alpha^{+W_\alpha} - \mathrm{Sat}(\mu)$, every curve $y \in Y_\mu^0$ must intersect the subsurface $U_\mu \subseteq S^z - \mu$ by part 4 of Lemma 5.8. Thus, subsurface projection gives a coarse map $\pi_{S^z - \mu} \colon Y_\mu \to \mathcal{C}'(S^z - \mu) = H(\mu)$. This map is uniformly coarsely Lipschitz by Lemmas 2.4 and 2.2. Since $\pi_{S^z - \mu}$ is also the identity on the vertices of $H(\mu)$, Lemma 2.3 ensures that the inclusion of $H(\mu)$ into Y_μ is a uniform quasi-isometric embedding. □

5.3.2 Tree guided type links

We now handle the case when μ is not of subsurface type. In this case, $\mathrm{lk}(\mu)$ is not all of $\mathcal{C}(S^z - \mu)$ as there are curves on S^z that are disjoint from μ but not in X_α, because their image under Π intersects α. This means we cannot rely on the hyperbolicity of $\mathcal{C}'(S^z - \mu)$ or use the subsurface projection map. Instead, we prove that $X_\alpha - \mathrm{Sat}(\mu)$ is F_i-guided for each fiber F_i where $\mu \cap F_i = \varnothing$. This implies Y_μ is a quasi-tree. We then show that $H(\mu)$ is quasi-isometrically embedded by using a combination of the subsurface projection maps and the tree guided structure of $X_\alpha - \mathrm{Sat}(\mu)$.

Proposition 5.17. *Let μ be a tree guided type simplex of X_α. For each F_i, if $\mu \cap F_i = \varnothing$, then $\mathrm{Sat}(\mu) \cap F_i = \varnothing$ and $X_\alpha - \mathrm{Sat}(\mu)$ is F_i-guided. Further, Y_μ is uniformly quasi-isometric to the uniform quasi-tree $X_\alpha - \mathrm{Sat}(\mu)$.*

Proof. Let F_i be a fiber so that $\mu \cap F_i = \varnothing$. We first show that $\mathrm{Sat}(\mu) \cap F_i$ must also be empty. Let c be a curve in $\mathrm{Sat}(\mu)$. By part 4 of Lemma 5.8, c is either a curve on $S^z - U_\mu$ or a curve of ∂U_μ that is not a core curve for an annular component of U_μ. This means $\Pi(c)$ can not be contained in $\Pi(U_\mu)$. However, we have $\alpha_i \subseteq \Pi(U_\mu)$ because $\mu \cap F_i = \varnothing$. Thus, $c \notin F_i$ when $c \in \mathrm{Sat}(\mu)$.

Since $\mathrm{Sat}(\mu) \cap F_i = \varnothing$, Lemma 4.5 says $X_\alpha - \mathrm{Sat}(\mu)$ is also F_i-guided. Lemma 4.3 then implies $X_\alpha - \mathrm{Sat}(\mu)$ is a uniform quasi-tree. To see that Y_μ is uniformly quasi-isometric to $X_\alpha - \mathrm{Sat}(\mu)$, we will show that for every $x, y \in Y_\mu^0$ that are not joined by an edge in

$X_\alpha - \mathrm{Sat}(\mu)$ but are joined by an edge in Y_μ the distance between x and y in $X_\alpha - \mathrm{Sat}(\mu)$ is at most 2.

The definition of edges in W_α ensures that there exists a multicurve ν so that $\nu \cup x$ and $\nu \cup y$ are adjacent vertices of W_α. Since $\Pi(\nu \cup x)$ is a pant decomposition of S that contains α, there exists $c \in \nu$ that is a curve in F_i. Since $\mathrm{Sat}(\mu) \cap F_i = \varnothing$, the curve c is a vertex of $X_\alpha - \mathrm{Sat}(\mu)$. Hence x and y are 2 apart in $X_\alpha - \mathrm{Sat}(\mu)$ as they are both disjoint from c. □

Proposition 5.18. *Let μ be a tree guided type simplex of X_α. The inclusion of $H(\mu)$ into the quasi-tree Y_μ is a $(15, 0)$-quasi-isometric embedding, making $H(\mu)$ a uniform quasi-tree.*

Proof. By the first item of Lemma 5.8, $\mathrm{lk}(\mu)$ is a join whenever U_μ is disconnected. Since $\mathrm{diam}(H(\mu)) \leqslant 2$ whenever $\mathrm{lk}(\mu)$ is a join, we can assume U_μ is connected. Further, we can assume that U_μ is not a single annulus as $\mathrm{lk}(\mu)$ would be a single vertex in that case. Since μ is not of subsurface type, Lemma 5.7 says U_μ must contain the marked point z as well as an infinite number of curve in $\Pi^{-1}(\alpha)$.

We will construct a coarsely Lipschitz map $\psi_\mu \colon Y_\mu \to H(\mu)$. Since each curve in Y_μ^0 must intersect U_μ, $\pi_{U_\mu}(y)$ is non-empty for each $y \in Y_\mu^0$. However, because $\mu \cap F_i = \varnothing$ it might be the case that none of the curve of $\pi_{U_\mu}(y)$ are vertices of X_α, that is, the curves in $\Pi(\pi_{U_\mu}(y))$ could intersect α on S. This prevents us from using the subsurface projection map alone to construct ψ_μ; instead we will use a combination of the subsurface projection map and the tree-guided structure of $X_\alpha - \mathrm{Sat}(\mu)$.

If the fiber F_i has $\mu \cap F_i = \varnothing$, then Proposition 5.17 implies that $\mathrm{Sat}(\mu) \cap F_i = \varnothing$ and Lemma 5.2 showed that $\mathrm{lk}(\mu) \cap F_i \subseteq Y_\mu$ is connected (in fact, it is the subtree of curves in F_i that are contained in U_μ). Because $\mathrm{lk}(\mu) \cap F_i$ is a connected subset of the tree F_i, the closest point projection of F_i onto $\mathrm{lk}(\mu) \cap F_i$ is well defined and Lipschitz. We will use this map to build a coarsely Lipschitz map from Y_μ to $H(\mu)$. The first step is to project each $y \in Y_\mu - H(\mu)$ with $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$ onto $\mathrm{lk}(\mu) \cap F_i$.

Let $y \in Y_\mu^0 - H^0(\mu)$ and suppose y is not disjoint from any curve in $\mathrm{lk}(\mu) \cap F_i$ where $\mu \cap F_i = \varnothing$, i.e., $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$. Let $y' = y$ if $y \in F_i$ and y' be any element of $\mathrm{lk}(y) \cap F_i$ if $y \notin F_i$. By Lemma 5.2, $\mathrm{lk}(y) \cap F_i$ is a connected subset of F_i in the latter case. Define $\psi_i(y)$ to be the closest point projection in F_i of y' onto $\mathrm{lk}(\mu) \cap F_i$. This is does not depend on the choice of y' because $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$ means $\mathrm{lk}(y) \cap F_i$ is a subtree of F_i that is disjoint from the subtree $\mathrm{lk}(\mu) \cap F_i$ when $y \notin F_i$.

We now verify that for each $y \in Y_\mu^0 - H^0(\mu)$ the set $\{\psi_i(y) : \mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing\}$ is a diameter 1 subset of $H(\mu)$ when it is non-empty.

Claim 5.19. Suppose $y \in Y_\mu^0 - H^0(\mu)$ so that $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$ and $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_j = \varnothing$ where $i \neq j$ and $\mu \cap F_i = \varnothing$ and $\mu \cap F_j = \varnothing$. Then $\psi_i(y)$ is connected by an X_α-edge to $\psi_j(y)$.

Proof. Because F_i and F_j are both contained in Y_μ and because $\Pi(y)$ will be disjoint from or equal to each of α_i and α_j on S, there exists $a_i \in F_i$ and $a_j \in F_j$ so that a_i and a_j are disjoint and $i(a_i, y) = 0$ and $i(a_j, y) = 0$; see Figure 4a. By the definition of ψ_i and ψ_j, we have $\psi_i(y) = \psi_i(a_i)$ and $\psi_j(y) = \psi_j(a_j)$.

Since $\mu \cap F_i$ and $\mu \cap F_j$ are both empty, there must exist a curve $b_j \in F_j$ so that b_j is disjoint from both μ and $\psi_i(a_i)$; see Figure 4b. Similarly, there must exists $b_i \in F_i$ so that

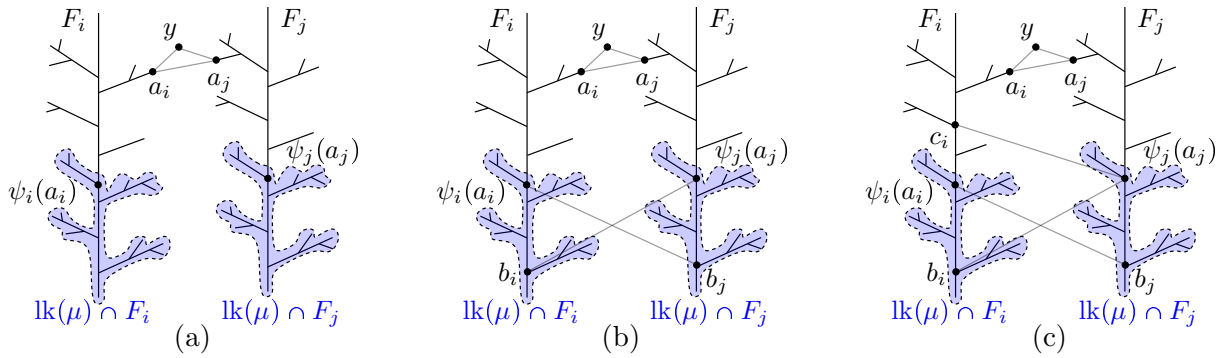

Figure 4: Schematic for the proof of Claim 5.19

b_i is disjoint from both μ and $\psi_j(a_j)$. Let γ_i be the geodesic in the tree F_i that connects a_i to $\psi_i(a_i)$. Because $\psi_i(a_i)$ is connected by an X_α-edge to b_j and $X_\alpha - \mathrm{Sat}(\mu)$ is F_j-guided, the union of the F_j-sequence (Definition 4.2) for γ_i is a connected subset of F_j that contains both a_j and b_j. Hence, the F_j-sequence of γ_i contains $\psi_j(a_j)$. There is therefore some vertex c_i of γ_i that is joined by an X_α-edge to $\psi_j(a_j)$; see Figure 4c. Since both c_i and b_i are in the subtree $\mathrm{lk}(\psi_j(a_j)) \cap F_i$, the F_i-geodesic connecting c_i and b_i is also contained in $\mathrm{lk}(\psi_j(a_j)) \cap F_i$ by Lemma 5.2. However, because b_i is in $\mathrm{lk}(\mu) \cap F_i$, the F_i-geodesic connecting c_i and b_i must pass through $\psi_i(a_i)$. This implies $\psi_i(a_i) \in \mathrm{lk}(\psi_j(a_j)) \cap F_i$. ■

We now use the ψ_i's and subsurface projection to define a coarse map $\psi_\mu \colon Y^0_\mu \to H^0(\mu)$:

- For $y \in H^0(\mu)$, define $\psi_\mu(y) = y$. In this case, y is a curve on U_μ, so $\psi_\mu(y) = \pi_{U_\mu}(y)$.
- For $y \in Y^0_\mu - H^0(\mu)$ we have two subcases:
  1. For all the fibers F_i where $\mu \cap F_i = \varnothing$, we have $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i \neq \varnothing$. Thus, for each fiber with $\mu \cap F_i = \varnothing$, there is $a_i \in \mathrm{lk}(\mu) \cap F_i$ that is disjoint from y. Since $a_i \subseteq U_\mu$ for every such a_i, $\pi_{U_\mu}(y)$ is disjoint from each a_i as well. Thus, $\pi_{U_\mu}(y)$ is a collection of curves in X_α that is contained in $\mathrm{lk}(\mu)$ and we define $\psi_\mu(y) = \pi_{U_\mu}(y)$. This is a uniformly bounded subset by Lemma 2.4.
  2. There exists a fiber F_i so that $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$, i.e., y is not disjoint from any curve in $\mathrm{lk}(\mu) \cap F_i$. In this case, we define

$$\psi_\mu(y) = \{\psi_i(y) : \mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing\}.$$

     By Claim 5.19, $\psi_\mu(y)$ is a diameter 1 subset of $H(\mu)$.

This definition means we can break $\psi_\mu(y)$ into two mutually exclusive case: one where $\psi_\mu(y) = \pi_{U_\mu}(y)$ and one where $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$ for some fiber F_i. In both case, $\psi_\mu(y)$ is uniformly bounded by either Lemma 2.4 or Claim 5.19.

Let y_1, y_2 be vertices of Y_μ that are connected by an edge of Y_μ. We will show that the $H(\mu)$-distance between $\psi_\mu(y_1)$ and $\psi_\mu(y_2)$ is uniformly bounded. This implies ψ_μ is coarsely Lipschitz (Lemma 2.2), which finishes the proof by Lemma 2.3. We first handle the case when y_1 and y_2 are actually disjoint curves, i.e., joined by an X_α-edge.

Claim 5.20. If $y_1, y_2 \in Y_\mu^0$ are disjoint curves, then $d_{H(\mu)}(\psi_\mu(y_1), \psi_\mu(y_2)) \leqslant 2$.

Proof. Let $y_1, y_2 \in Y_\mu^0$ be disjoint curves.

Assume first that $\psi_\mu(y_1) = \pi_{U_\mu}(y_1)$ and $\psi_\mu(y_2) = \pi_{U_\mu}(y_2)$. Without loss of generality, we can assume $y_1 \notin H(\mu)$. If $y_2 \in H(\mu)$, then $y_2 = \psi_\mu(y_2) = \pi_{U_\mu}(y_2)$ does not intersect $\psi_\mu(y_1) = \pi_{U_\mu}(y_1)$ by Lemma 2.4. If instead $y_2 \notin H(\mu)$, then $\mathrm{lk}(\mu) \cap F_i$ contains a curve of both $\mathrm{lk}(y_1) \cap F_i$ and $\mathrm{lk}(y_2) \cap F_i$ for each fiber with $F_i \cap \mu = \varnothing$. Further, $\mathrm{lk}(y_1) \cap \mathrm{lk}(y_2) \cap F_i$ is non-empty and connected because y_1 and y_2 are joined by an edge of $X_\alpha - \mathrm{Sat}(\mu)$ and $X_\alpha - \mathrm{Sat}(\mu)$ is F_i-guided. Therefore, the Helly property of trees implies that $\mathrm{lk}(y_1) \cap \mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_i \neq \varnothing$. Thus, there is a curve $c \in X_\alpha - \mathrm{Sat}(\mu)$ that is contained in U_μ and disjoint from both y_1 and y_2. This implies $d_{H(\mu)}(\psi_\mu(y_1), \psi_\mu(y_2)) \leqslant 2$ because both $\psi_\mu(y_1) = \pi_{U_\mu}(y_1)$ and $\psi_\mu(y_2) = \pi_{U_\mu}(y_2)$ will not intersect c.

Next assume $\mathrm{lk}(y_1) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$ and $\mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$ for single fiber F_i. There must exist $a \in F_i$ so that $i(a, y_1) = 0$ and $i(a, y_2) = 0$. Thus, $d_{H(\mu)}(\psi_\mu(y_1), \psi_\mu(y_2)) = 0$ because $\psi_i(y_1) = \psi_i(a) = \psi_i(y_2)$.

Now assume $\mathrm{lk}(y_1) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$ and $\mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_j = \varnothing$ for distinct fibers F_i and F_j with $\mu \cap F_i$ and $\mu \cap F_j$ both empty. There exist $a_i \in F_i$ and $a_j \in F_j$ that are both disjoint from or equal to each of y_1 and y_2 (however a_i and a_j might not be disjoint). We can now employ a similar argument to the proof of Claim 5.19. There must be curves $b_i \in F_i$ and $b_j \in F_j$ that are disjoint from μ and respectively disjoint from $\psi_j(y_2) = \psi_j(a_j)$ and $\psi_i(y_1) = \psi_i(a_i)$. Let γ_i be the concatenation of a F_i-geodesic connecting $\psi_i(a_i)$ to a_i with the X_α-edge from a_i to y_1. By examining the F_j-sequence for γ_i, we find a vertex c_i on the F_i-geodesic between a_i and $\psi_i(a_i)$ that is disjoint from $\psi_j(a_j)$. As in the end of Claim 5.19, this forces $\psi_j(a_j) = \psi_j(y_2)$ to be joined by an X_α-edge to $\psi_i(a_i) = \psi_i(y_1)$ because $\mathrm{lk}(\psi_j(a_j)) \cap F_i$ is connected. Since $\psi_i(y_1) \subseteq \psi_\mu(y_1)$ and $\psi_j(y_2) \subseteq \psi_\mu(y_2)$, this shows that the $H(\mu)$-distance between $\psi_\mu(y_1)$ and $\psi_\mu(y_2)$ is at most 1.

Finally, assume $\mathrm{lk}(y_1) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$, but $\psi_\mu(y_2) = \pi_{U_\mu}(y_2)$. In this case, y_2 must be disjoint from a curve in $\mathrm{lk}(\mu) \cap F_i$, i.e., $\mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_i \neq \varnothing$. In particular, $\psi_\mu(y_2) = \pi_{U_\mu}(y_2)$ does not intersect an element of $\mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_i$. Thus, we can prove $d_{H(\mu)}(\psi_\mu(y_1), \psi_\mu(y_2)) \leqslant 1$ by proving that $\psi_i(y_1) \subseteq \psi_\mu(y_1)$ is an element of $\mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_i$.

Since y_1 and y_2 are disjoint, there exists $a \in F_i$ so that $i(a, y_1) = 0$ and $i(a, y_2) = 0$. Since $a \in \mathrm{lk}(y_2) \cap F_i$ and $\mathrm{lk}(y_2) \cap F_i$ is a subtree of F_i (Lemma 5.2), the geodesic in F_i from a to any point in $\mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_i$ is contained in $\mathrm{lk}(y_2) \cap F_i$. By construction $\psi_i(y_1)$ is the endpoint of the unique geodesic in F_i connecting a to $\mathrm{lk}(\mu) \cap F_i$. Since this path is a subpath of any geodesic in F_i that connects a to a point in $\mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_i$, we have that $\psi_i(y_1) \in \mathrm{lk}(y_2) \cap \mathrm{lk}(\mu) \cap F_i$ as desired. ∎

To complete the proof that ψ_μ is coarsely Lipschitz, suppose $y_1, y_2 \in Y_\mu^0$ are Y_μ-adjacent but not X_α-adjacent. There then exists some multicurve ν so that $\nu \cup y_1$ and $\nu \cup y_2$ are adjacent pants decompositions in W_α. Since ν is a single curve away from being a maximal simplex of X_α, ν must contain a curve from each fiber F_i. Because $\mathrm{Sat}(\mu) \cap F_i = \varnothing$ for some fiber F_i, $Y_\mu = X_\alpha^{+W_\alpha} - \mathrm{Sat}(\mu)$ contains all of this fiber F_i. Hence, the multicurve ν must contain a curve c that is a vertex of Y_μ. Thus there is $c \in Y_\mu$ with c disjoint from both y_1 and y_2. Therefore, Claim 5.20 implies

$$d_{H(\mu)}(\psi_\mu(y_1), \psi_\mu(y_2)) \leqslant d_{H(\mu)}(\psi_\mu(y_1), \psi_\mu(c)) + d_{H(\mu)}(\psi_\mu(c), \psi_\mu(y_2)) + 1 \leqslant 5.$$

Lemma 2.2 now implies ψ_μ is $(15, 0)$-coarsely Lipschitz. □

Remark 5.21 (ψ_μ versus closest point projection)**.** Because $H(\mu)$ is quasi-isometrically embedded in Y_μ and Y_μ is hyperbolic there is a coarse closest point projection $\mathfrak{p}_\mu \colon Y_\mu \to H(\mu)$ defined by sending y to the uniformly bounded set $\{x \in H(\mu) : d_{Y_\mu}(x, y) = d_{Y_\mu}(x, H(\mu))\}$. This closest point projection is within uniformly bounded distance of the map ψ_μ defined in the proof of the above theorem. To see this, let $y \in Y^0_\mu$. If $\psi_\mu(y) = \pi_{U_\mu}(y)$, then y and $\psi_\mu(y)$ are at most 2 apart by Lemma 2.4. If instead $\mathrm{lk}(y) \cap \mathrm{lk}(\mu) \cap F_i = \varnothing$ for some fiber with $\mu \cap F_i = \varnothing$, then let γ be the Y_μ-geodesic connecting y to $\mathfrak{p}_\mu(y)$. Let z be the endpoint of γ in $\mathfrak{p}_\mu(y)$. There must exist a curve $b \in F_i \cap \mathrm{lk}(\mu)$ so that either $b = z$ or z and b are disjoint. Thus, as in the proof of Claim 5.19, examining the F_i-sequence of γ produces vertex c of γ that is distance 1 from $\psi_i(y)$. Since $\psi_i(y) \in H(\mu)$ and z minimizes the distance from y to $H(\mu)$, c and z must be at most 1 apart. Thus, $\psi_\mu(y)$ within distance 2 from $\mathfrak{p}_\mu(y)$.

5.4 (X_α, W_α) is a combinatorial HHS.

We now combine the result of the previous two sections to conclude that (X_α, W_α) is a combinatorial HHS with a cobounded action by E_α

Theorem 5.22. *The graph W_α is connected and a hierarchically hyperbolic space because the pair (X_α, W_α) is a combinatorial HHS. Further, the action of E_α on S^z induces a cobounded action of E_α on W_α.*

Proof. We first check the five parts of Definition 2.15 for some $\delta \geqslant 2$ ultimately depending only on S. Theorem 2.18 will then imply that W is connected and a hierarchically hyperbolic space.

1. If $\mu_1, \ldots, \mu_n$ is a sequence of non-maximal simplices of X_α so that

   $$\mathrm{lk}(\mu_1) \subsetneq \mathrm{lk}(\mu_2) \subsetneq \cdots \subsetneq \mathrm{lk}(\mu_n),$$

   then Lemma 5.8 implies n is bounded in terms of the complexity of S.

2. Corollary 5.13 proves that $X_\alpha^{+W_\alpha}$ is quasi-isometric to a tree and hence hyperbolic.

3. The links of simplices of X_α split into subsurface and tree-guided types; see Definition 5.14. The proof of uniform hyperbolicity and uniform quasi-isometric embedding for subsurface type is shown in Propositions 5.15 and 5.16, while the proof for the tree-guided type is completed in Proposition 5.18.

4. Because $\mathrm{diam}(H(\mu)) > 2$ implies $\mathrm{lk}(\mu)$ is not a join or a single vertex, this follows from Proposition 5.9.

5. This condition is ensured by Proposition 5.12.

To see that the action of E_α on W_α is cobounded, start with the fact that Π will map each vertex of W_α to a pants decomposition of S containing the curves in α. Recall, E_α fits into the short exact sequence

$$1 \to \mathrm{Push}(S, z) \to E_\alpha \to G_\alpha \to 1$$

and Π is equivariant with respect to the quotient map $E_\alpha \to G_\alpha$. For each pants decomposition ρ of S with $\alpha \subseteq \rho$, there only exists finitely many $\mathrm{Push}(S;z)$-orbits of elements of $\{\mu \in W^0_\alpha : \Pi(\mu) = \rho\}$. Moreover, there exist only finitely many $G_\alpha = \mathrm{Stab}_{\mathrm{MCG}(S)}(\alpha)$-orbits of pants decompositions of $S - \alpha$. Hence, there are only finitely many E_α-orbits of vertices of W_α, making the E_α-action cobounded. □

6 Adding in annuli

We continue with the assumptions of Section 5. That is, $S \cong S^0_{g,0}$ with $g \geqslant 2$, and $\alpha = \alpha_1 \cup \cdots \cup \alpha_m$ is a fixed multicurve on S. For a fixed $z \in S$, the group E_α is then the full preimage in $\mathrm{MCG}(S;z)$ of the stabilizer of α in $\mathrm{MCG}(S)$. The spaces X_α and W_α are as defined in Definitions 5.1 and 5.5.

The extension group E_α does not act metrically properly on the space W_α. This is because the action of E_α on a maximal simplex in X_α has large stabilizer, namely the Dehn twists around the curves in the simplex. To fix this, we need to "blow-up" X_α by adding marking arcs for each of the curves in X_α that account for the action of Dehn twists. We denote this blow up $B(X_\alpha)$. The corresponding $B(X_\alpha)$-graph will be denoted $B(W_\alpha)$. As in the previous section, we will as a quantity is *uniform* if it does not depend on the multicurve α, the surface S, or specific simplices/vertices of $B(X_\alpha)$ and $B(W_\alpha)$.

6.1 Defining the blow-ups of X_α and W_α

For a curve c in S^z, let A_c be the annulus with core curve c. For each $c \in X^0_\alpha$, let $B(c)$ denote the set of vertices of the annular complex $\mathcal{C}(A_c)$. We call the elements of $B(c)$ the *marking arcs* for c. Let $B(X_\alpha)$ be the space obtained from X_α by adding in the following vertices and edges:

1. For each curve $c \in X_\alpha$, add the elements of $B(c)$ as additional vertices.
2. For each curve $c \in X_\alpha$, add an edge between c and each element of $B(c)$.
3. If c and c' are disjoint curves in X_α, then for each $v \in B(c)$ and $v' \in B(c') \cup \{c'\}$ add an edge connecting v and v'.

For a simplex μ of $B(X_\alpha)$ we define three sets of vertices of $B(X_\alpha)$:

- $\mathrm{base}(\mu)$ is the set $\{c \in X^0_\alpha : c \in \mu\}$.
- $\mathrm{supp}(\mu)$ is the set $\mathrm{base}(\mu) \cup \{c \in X^0_\alpha : B(c) \cap \mu \neq \varnothing\}$.
- $\mathrm{markarc}(\mu)$ is the set of vertices of $\mu - \mathrm{base}(\mu)$.

We call the curves in $\mathrm{base}(\mu)$ and $\mathrm{supp}(\mu)$ the *base* and *support* curves of μ respectively. The vertices of $\mathrm{markarc}(\mu)$ are called the *marking arcs* of μ. Note, each curve of $\mathrm{supp}(\mu) - \mathrm{base}(\mu)$ corresponds to a marking arc of μ that does not have a base curve in μ. We say that $c \in \mathrm{base}(\mu)$ is an *unmarked base curve of* μ if $B(c) \cap \mathrm{markarc}(\mu) = \varnothing$. We denote the set

of unmarked base curves of μ by unmark(μ). The maximal simplices of $B(X_\alpha)$ are precisely the (not necessarily clean) markings of S^z whose base curves are maximal simplices in X_α.

We now define our $B(X_\alpha)$-graph which we denote $B(W_\alpha)$. Recall, for $t, t' \in B(c) = \mathcal{C}^0(A_c)$, $d_{A_c}(t, t')$ denotes the distance between t and t' in the annular complex $\mathcal{C}(A_c)$. The graph $B(W_\alpha)$ has vertex set the maximal simplices of $B(X_\alpha)$ with edges between two vertices $\mu, \nu \in B(W_\alpha)$ corresponding to *twist* and *flip* moves:

- (twist) base(μ) = base(ν) and there exists exactly one curve $c \in$ base(μ) so that markarc(μ) $\cap B(c) \neq$ markarc(ν) $\cap B(c)$. Further, for this curve c, if $t \in$ markarc(μ) $\cap B(c)$ and $t' \in$ markarc(ν) $\cap B(c)$, then $d_{A_c}(t, t') = 1$.
- (flip) There exists curves $c \in$ base(μ) and $d \in$ base(ν) so that
  - (base(μ) $- c$) $\cup d =$ base(ν) and $i(c, d) \leqslant 4$, i.e., base(μ) and base(ν) differ by an edge in W_α;
  - if $t_c \in$ markarc(μ) and $t_d \in$ markarc(ν) are the marking arcs for c and d, then
$$d_{A_c}(t_c, d) \leqslant 1 \text{ and } d_{A_d}(t_d, c) \leqslant 1;$$
  - for all curves $b \in$ base(μ) $- c =$ base(ν) $- d$, the marking arc of b in μ is the same as the marking arc of b in ν.

The graphs $B(X_\alpha)$ and $B(W_\alpha)$ are related to X_α and W_α by simplicial maps induced by sending every simplex $\mu \subseteq B(X_\alpha)$ to supp(μ). This will allow us to reduce properties of $(B(X_\alpha), B(W_\alpha))$ to properties about (X_α, W_α).

Lemma 6.1. *The following maps on vertices extend to simplicial maps between the given complexes.*

- $B^0(X_\alpha) \to X_\alpha$ *by* $x \to$ supp(x).
- $B^0(X_\alpha)^{+B(W_\alpha)} \to X_\alpha^{+W_\alpha}$ *by* $x \to$ supp(x).
- $B^0(W_\alpha) \to W_\alpha$ *by* $\mu \to$ supp(μ).

Proof. In all three cases, it suffices to show that two vertices joined by an edge in the domain graph have supports that are either equal or joined by an edge in the codomain graph.

Let x, y be vertices of $B(X_\alpha)$ that are joined by an edge. Thus either supp(x) = supp(y) or supp(x) is disjoint from supp(y). Hence supp(x) and supp(y) are either equal or joined by an edge in X_α.

Let x, y be vertices of $B(X_\alpha)^{+B(W_\alpha)}$ that are joined by an edge. By the previous paragraph, we can assume x and y are joined by a $B(W_\alpha)$-edge. If x and y are joined by an edge coming from a twist move in $B(W_\alpha)$, then supp(x) = supp(y). If instead x and y are joined be an edge coming from a flip move, then supp(x) and supp(y) fill a connected complexity 1 subsurface of S^z and intersect at most 4 times. Hence, supp(x) and supp(y) will be joined by a W_α-edge of $X_\alpha^{W_\alpha}$.

Let μ, ν be vertices of $B(W_\alpha)$ that are joined by an edge. If μ and ν are joined by a twist move, then supp(μ) = supp(ν). If μ and ν are joined by a flip move, then supp(μ) and supp(ν) are joined by a flip move in W_α by definition. □

6.2 The combinatorial conditions with annuli

We now verify that $(B(X_\alpha), B(W_\alpha))$ satisfies the combinatorial condition of being a combinatorial HHS. We begin by describing the links of simplices of $B(X_\alpha)$. Henceforth, we use $\mathrm{lk}_B(\cdot)$ to denote the link of simplices in $B(X_\alpha)$ and $\mathrm{lk}_X(\cdot)$ to denote the link of simplices in X_α.

Lemma 6.2. *For any non-maximal simplex μ of $B(X_\alpha)$, the link of μ in $B(X_\alpha)$ is the subgraph spanned by the union of the following vertices:*

1. *The vertices of $B(c)$ for each $c \in \mathrm{unmark}(\mu)$.*
2. *The set of vertices $\{x \in B^0(X_\alpha) : \mathrm{supp}(x) \in \mathrm{lk}_X(\mathrm{supp}(\mu))\}$.*
3. *The vertices of $\mathrm{supp}(\mu) - \mathrm{base}(\mu)$.*

Moreover, if $\mathrm{supp}(\mu) \neq \mathrm{base}(\mu)$, then $\mathrm{lk}_B(\mu)$ is either a single vertex (which will be $\mathrm{supp}(\mu) - \mathrm{base}(\mu)$) or is a join of some graph with $\mathrm{supp}(\mu) - \mathrm{base}(\mu)$.

Proof. If $c \in \mathrm{unmark}(\mu)$, then all the vertices of $B(c)$ are in $\mathrm{lk}_B(\mu)$ by the definition of the edges in $B(X_\alpha)$.

If $x \in B^0(X_\alpha)$ with $\mathrm{supp}(x) \in \mathrm{lk}_X(\mathrm{supp}(\mu))$, then $\mathrm{supp}(x)$ is disjoint from all the curves in $\mathrm{supp}(\mu)$. By the edge relations in $B(X_\alpha)$, this makes $x \in \mathrm{lk}_B(\mu)$.

Finally, suppose c is a vertex of $\mathrm{supp}(\mu) - \mathrm{base}(\mu)$. Since μ is a simplex, c must be disjoint from $\mathrm{base}(\mu)$. However, this means c must be joined by an edge of every element of μ by the definition of the $B(X_\alpha)$ edges.

To verify these are the only elements of $\mathrm{lk}_B(\mu)$ let x be a vertex of $\mathrm{lk}_B(\mu)$. If $x \in X^0_\alpha$, then either $x \in \mathrm{lk}_X(\mathrm{supp}(\mu))$ or $x \in \mathrm{supp}(\mu) - \mathrm{base}(\mu)$. If instead x is in $B(c)$ for some curve $c \in X^0_\alpha$, then either $c \in \mathrm{unmark}(\mu)$ or $c \in \mathrm{lk}_X(\mathrm{supp}(\mu))$.

For the moreover statement, let $c \in \mathrm{supp}(\mu) - \mathrm{base}(\mu)$ and let t_c be the element of $\mathrm{markarc}(\mu) \cap B(c)$. Any vertex $x \in \mathrm{lk}_B(\mu)$ will be joined by an edge to t_c by definition. However, this makes x and c joined by an edge from the the definition of edges in $B(X_\alpha)$. □

We now describe the nesting of links of $B(X_\alpha)$ by using the topology of subsurfaces of S^z. This is an extension of Lemma 5.8 from X_α to $B(X_\alpha)$ and is our main tool in this section. Recall that Definition 5.6 defined $U_{\mathrm{supp}(\mu)}$ as subsurface filled by the curves in $\mathrm{lk}_X(\mathrm{supp}(\mu))$. We remark that it is possible that $U_{\mathrm{supp}(\mu)}$ is empty for a non-maximal simplex $\mu \subseteq B(X_\alpha)$. This occurs when $\mathrm{supp}(\mu)$ is a pants decomposition of S^z, but not every base curve of μ has a marking arc.

Lemma 6.3. *Suppose μ and μ' are non-maximal simplices of $B(X_\alpha)$ with $\mathrm{supp}(\mu) = \mathrm{base}(\mu)$ and $\mathrm{supp}(\mu') = \mathrm{base}(\mu')$. We have $\mathrm{lk}_B(\mu) \subseteq \mathrm{lk}_B(\mu')$ if and only if*

- *$U_{\mathrm{supp}(\mu)} \subseteq U_{\mathrm{supp}(\mu')}$; and*
- *for each $c \in \mathrm{unmark}(\mu)$, either $c \in \mathrm{unmark}(\mu')$ or $c \subseteq U_{\mathrm{supp}(\mu')}$.*

Remark 6.4. The above conditions for the nesting of links can be stated purely in terms of subsurfaces by using annuli. For a simplex μ where $\mathrm{supp}(\mu) = \mathrm{base}(\mu)$ define C_μ to be the disjoint union of the annuli $\{A_c : c \in \mathrm{unmark}(\mu)\}$. The conclusion of Lemma 6.3 can then be phrased as $\mathrm{lk}_B(\mu) \subseteq \mathrm{lk}_B(\mu')$ if and only if $U_{\mathrm{supp}(\mu)} \subseteq U_{\mathrm{supp}(\mu')}$ and $C_\mu \subseteq U_{\mathrm{supp}(\mu')} \cup C_{\mu'}$.

Proof of Lemma 6.3. Before starting the proof, observe that by applying Lemma 6.2, the assumption that $\text{base}(\mu) = \text{supp}(\mu)$ implies that a vertex $x \in \text{lk}_B(\mu)$ if and only if $\text{supp}(x) \in \text{lk}_X(\text{supp}(\mu))$ or $x \in B(c)$ for some $c \in \text{unmark}(\mu)$.

Suppose $\text{lk}_B(\mu) \subseteq \text{lk}_B(\mu')$. For the first bullet, we prove $\text{lk}_X(\text{supp}(\mu)) \subseteq \text{lk}_X(\text{supp}(\mu'))$, which implies $U_{\text{supp}(\mu)} \subseteq U_{\text{supp}(\mu')}$ by Lemma 5.8. If $x \in \text{lk}_X(\text{supp}(\mu))$, then $x \in \text{lk}_B(\mu)$ and $x \in X_\alpha$. Since $\text{lk}_B(\mu) \subseteq \text{lk}_B(\mu')$, x must be an element of $\text{lk}_B(\mu')$ that is also an element of X_α. By the above observation, this only occurs if $x = \text{supp}(x) \in \text{lk}_X(\text{supp}(\mu'))$ as well. This shows $\text{lk}_X(\text{supp}(\mu)) \subseteq \text{lk}_X(\text{supp}(\mu'))$ as desired.

For the second bullet, let c be an unmarked base curve of μ. By Lemma 6.2, this implies $B(c) \subseteq \text{lk}_B(\mu)$, which means $B(c)$ is also contained in $\text{lk}_B(\mu')$. Thus either $c \in \text{unmark}(\mu')$ or $c \in \text{lk}_X(\text{supp}(\mu'))$. In the latter cases, we have that $c \subseteq U_{\text{supp}(\mu')}$ by Definition 5.6.

Now suppose $U_{\text{supp}(\mu)} \subseteq U_{\text{supp}(\mu')}$ and for all $c \in \text{unmark}(\mu)$, either $c \subseteq U_{\text{supp}(\mu')}$ or $c \in \text{unmark}(\mu')$. Let $x \in \text{lk}_B(\mu)$. By the observation at the beginning, we know that either $\text{supp}(x) \in \text{lk}_X(\text{supp}(\mu))$ or x is a marking arc for some unmarked base curve c of μ. In the first case, Lemma 5.8 implies $\text{supp}(x) \in \text{lk}_X(\text{supp}(\mu'))$ because $U_{\text{supp}(\mu)} \subseteq U_{\text{supp}(\mu')}$. Thus, $x \in \text{lk}_B(\mu')$ by Lemma 6.2. In the second case, we either have $c \in \text{unmark}(\mu')$ or $c \subseteq U_{\text{supp}(\mu')}$ by hypothesis. In both cases $x \in \text{lk}_B(\mu')$ by Lemma 6.2. Thus we have $\text{lk}_B(\mu) \subseteq \text{lk}_B(\mu')$ □

Using Lemma 6.3, we show that $(B(X_\alpha), B(W_\alpha))$ satisfies the combinatorial conditions of being a combinatorial HHS. The essence of these arguments are the same as Proposition 5.9 and Proposition 5.12 in the non-annular case.

Proposition 6.5. *Suppose μ and μ' are non-maximal simplices of $B(X_\alpha)$ so that there exists a non-maximal simplex ν of X_α with $\text{lk}_B(\nu) \subseteq \text{lk}_B(\mu) \cap \text{lk}_B(\mu')$ and $\text{lk}_B(\nu)$ not a join or a single vertex. There exists a (possibly empty) simplex ρ of $\text{lk}_B(\mu')$ so that $\text{lk}_B(\mu' \bowtie \rho) \subseteq \text{lk}_B(\mu)$ and if ν is any simplex as in the proceeding sentence, then $\text{lk}_B(\nu) \subseteq \text{lk}_B(\mu' \bowtie \rho)$.*

Proof. Let $\eta = \text{markarc}(\mu) \bowtie \text{supp}(\mu)$ and $\eta' = \text{markarc}(\mu') \bowtie \text{supp}(\mu')$. Lemma 6.2 implies that $\text{lk}_B(\eta)$ is $\text{lk}_B(\mu)$ minus the vertices of $\text{supp}(\mu) - \text{base}(\mu)$ and that the same holds for η' and μ'. Thus, $\text{lk}_B(\mu) = \text{lk}_B(\eta) \bowtie (\text{supp}(\mu) - \text{base}(\mu))$ and the same holds for η' and μ'. We first verify that it suffices to prove the proposition is true for η and η', instead of μ and μ'. The advantage of working with η and η' over μ and μ' is that Lemma 6.3 applies to η and η' because $\text{supp}(\eta) = \text{base}(\eta)$ and $\text{supp}(\eta') = \text{base}(\eta')$.

Claim 6.6. If ρ is a simplex that satisfies Proposition 6.5 for η and η', then the join of ρ and $\text{supp}(\mu') - \text{base}(\mu')$ is a simplex that satisfies the proposition for μ and μ'.

Proof. Let ρ be the simplex that satisfies Proposition 6.5 for η and η' and let $\rho' = \rho \bowtie (\text{supp}(\mu')) - \text{base}(\mu'))$. Since $\rho \subseteq \text{lk}_B(\eta')$ and $\text{lk}_B(\mu') = \text{lk}_B(\eta') \bowtie (\text{supp}(\mu')) - \text{base}(\mu'))$, ρ' is in fact a simplex of $B(X_\alpha)$. Moreover, $\eta' \cup \rho = \mu' \cup \rho'$, hence $\text{lk}_B(\eta' \bowtie \rho) = \text{lk}_B(\mu' \bowtie \rho')$. As ρ satisfies Proposition 6.5 for η and η', we have $\text{lk}_B(\eta' \bowtie \rho) = \text{lk}_B(\mu' \bowtie \rho') \subseteq \text{lk}_B(\eta)$. Since $\text{lk}_B(\eta) \subset \text{lk}_B(\mu)$, this gives $\text{lk}_B(\mu' \bowtie \rho') \subseteq \text{lk}_B(\mu)$ as desired.

We now show the second requirement of Proposition 6.5 for ρ' and μ', μ. If there exists a non-maximal simplex $\nu \subseteq B(X_\alpha)$, so that $\text{lk}_B(\nu) \subseteq \text{lk}_B(\mu)$ and $\text{lk}_B(\nu) \nsubseteq \text{lk}_B(\eta)$, then $\text{supp}(\mu) - \text{base}(\mu) \neq \varnothing$ and $\text{lk}_B(\nu)$ must contain vertices of $\text{supp}(\mu) - \text{base}(\mu)$. In this case, Lemma 6.2 says $\text{lk}_B(\mu)$ is either a single vertex contained in $\text{supp}(\mu) - \text{base}(\mu)$ or the join of some graph with the vertices in $\text{supp}(\mu) - \text{base}(\mu)$. Hence, $\text{lk}_B(\nu)$ would need to be either a

single vertex or a join as well. Since the same reasoning applies to $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\mu')$, we have that whenever $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\mu) \cap \mathrm{lk}_B(\mu')$ and $\mathrm{lk}_B(\nu)$ is not a join or a single vertex, then $\mathrm{lk}_B(\nu)$ must be contained in $\mathrm{lk}_B(\eta) \cap \mathrm{lk}_B(\eta')$. Since ρ satisfies Proposition 6.5 for η and η', then any such simplex ν will have $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\eta' \bowtie \rho) = \mathrm{lk}_B(\mu' \bowtie \rho')$. ■

We now construct the required simplex ρ for η and η'. We will define two simplices, ρ_U and ρ_A, of $B(X_\alpha)$ so that the desired ρ will be $\rho_U \bowtie \rho_A$.

To define ρ_A, let $\mathcal{A}$ to be the set of curves $c \in \mathrm{unmark}(\eta')$ so that $c \notin \mathrm{unmark}(\eta)$ and $c \nsubseteq U_{\mathrm{supp}(\eta)}$. If $\mathcal{A} = \varnothing$, then define $\rho_A = \varnothing$. Otherwise, for each $c \in \mathcal{A}$, pick any vertex $t_c \in B(c)$, then define $\rho_A = \{t_c : c \in \mathcal{A}\}$.

The definition of ρ_U will be similar to the definition of ρ in the proof of Proposition 5.9 for the non-annular case, with extra care to keep track of the set $B(c)$ that are contained in $\mathrm{lk}_B(\eta) \cap \mathrm{lk}_B(\eta')$. Let V_0 be the subsurface of S^z filled by the curves in $\mathrm{lk}_X(\mathrm{supp}(\eta)) \cap \mathrm{lk}_X(\mathrm{supp}(\eta'))$, then let V be the disjoint union of the non-annular components of V_0. Let σ be the set of curves in $\mathrm{unmark}(\eta)$ that are contained in $U_{\mathrm{supp}(\eta')}$. Since σ contains base curves of η, none of the curves of σ can be contained in V_0, thus $\sigma \subseteq U_{\mathrm{supp}(\eta')} - V_0$. The proof of Claim 5.10 implies ∂V is a simplex of X_α, hence Lemma 5.4 produces a pants decomposition τ of $U_{\mathrm{supp}(\eta')} - V_0$ so that $\sigma \subseteq \tau$ and $\tau \subseteq X_\alpha$. For each $c \in \tau \cup \partial V_0$, if $c \notin \mathrm{unmark}(\eta)$ and $c \nsubseteq U_{\mathrm{supp}(\eta)}$, select a marking arc $t_c \in B(c)$. Define

$$\rho_U = \tau \cup \partial V_0 \cup \{t_c : c \in \tau \cup \partial V_0 \text{ with } c \notin \mathrm{unmark}(\eta) \text{ and } c \nsubseteq U_{\mathrm{supp}(\eta)}\}.$$

We now verify that $\rho = \rho_A \bowtie \rho_U$ has the three desired properties given in the proposition. Note, $\mathrm{supp}(\eta' \bowtie \rho) = \mathrm{base}(\eta' \bowtie \rho) = \mathrm{base}(\eta') \cup \mathrm{base}(\rho_U)$ by construction.

- $\rho \subseteq \mathrm{lk}_B(\eta')$: This is immediate from the construction of ρ, as every vertex $x \in \rho$ is either a marking arc for an unmarked base curve of η' or has $\mathrm{supp}(x) \in \mathrm{lk}_X(\mathrm{supp}(\eta'))$.

- $\mathrm{lk}_B(\eta' \bowtie \rho) \subseteq \mathrm{lk}_B(\eta)$: Since $\mathrm{base}(\rho_U) = \tau \cup \partial V$, we have $U_{\mathrm{supp}(\eta' \bowtie \rho)} = U_{\mathrm{supp}(\eta') \bowtie (\tau \cup \partial V)}$. Thus, a repeat of the proof of Claim 5.11 proves that $U_{\mathrm{supp}(\eta' \bowtie \rho)} = U_{\mathrm{supp}(\eta') \bowtie (\tau \cup \partial V)} = V$. Since V is filled by a subset of curves in $\mathrm{lk}_X(\mathrm{supp}(\eta))$, we have $U_{\mathrm{supp}(\eta' \bowtie \rho)} = V \subseteq U_{\mathrm{supp}(\eta)}$. This verifies the first condition in Lemma 6.3.

  For the second condition in Lemma 6.3, let c be an unmarked base curve of $\eta' \bowtie \rho$. If $c \in \mathrm{base}(\eta')$, then the choice of ρ_A ensures that c is either an unmarked base curve of η or is contained on $U_{\mathrm{supp}(\eta)}$. If instead $c \in \mathrm{base}(\rho_U)$, then the choice of marking arcs in ρ_U again ensures that c is either an unmarked base curve of η or is contained in $U_{\mathrm{supp}(\eta)}$. Hence, $\mathrm{lk}_B(\eta' \bowtie \rho) \subseteq \mathrm{lk}_B(\eta)$.

- If ν is a non-maximal simplex of $B(X_\alpha)$ with $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\eta) \cap \mathrm{lk}_B(\eta')$, then either $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\eta' \bowtie \rho)$ or $\mathrm{lk}_B(\nu)$ is a single vertex or a join: Assume $\mathrm{lk}_B(\nu)$ is not a single vertex or a join and $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\eta) \cap \mathrm{lk}_B(\eta')$. Because $\mathrm{supp}(\nu) \neq \mathrm{base}(\nu)$ would imply that $\mathrm{lk}_B(\nu)$ is a join or a single vertex by Lemma 6.2, we know $\mathrm{supp}(\nu) = \mathrm{base}(\nu)$ and can apply Lemma 6.3 to ν.

  If ν contains an unmarked base curve c, then every vertex in $\mathrm{lk}_B(\nu) - B(c)$ will be joined by an edge to every vertex of $B(c)$. Since $\mathrm{lk}_B(\nu)$ is not a join, this means either $\mathrm{lk}_B(\nu) = B(c)$ or ν does not have any unmarked base curves.

If $\mathrm{lk}_B(\nu) = B(c)$, then $U_{\mathrm{supp}(\nu)} = \varnothing$ and c is the only unmarked base curve of ν. Since $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\eta) \cap \mathrm{lk}_B(\eta')$, we have four possibilities coming from Lemma 6.3. In each case we verify that c is either an unmarked base curve of $\eta' \bowtie \rho$ or is contained in $U_{\mathrm{supp}(\eta' \bowtie \rho)} = V$.

- If $c \in \mathrm{unmark}(\eta) \cap \mathrm{unmark}(\eta')$ or $c \subseteq U_{\mathrm{supp}(\eta)}$ and $c \in \mathrm{unmark}(\eta')$, then c is an unmarked base curve of $\eta' \bowtie \rho$, because of the choice of the marking arcs in ρ_A.
- If $c \in \mathrm{unmark}(\eta)$ and $c \subseteq U_{\mathrm{supp}(\eta')}$, then c is a curve in the multicurve σ defined while constructing ρ_U. Hence $c \in \mathrm{base}(\rho_U)$ and c is unmarked by the choice of marking arcs in ρ_U.
- If $c \subseteq U_{\mathrm{supp}(\eta)}$ and $c \subseteq U_{\mathrm{supp}(\eta')}$, then $c \subseteq V_0$ because $c \in \mathrm{lk}_X(\mathrm{supp}(\eta)) \cap \mathrm{lk}_X(\mathrm{supp}(\eta'))$. If c is a core curve of an annular component of V_0, then $c \in \partial V_0$. Thus $c \in \mathrm{base}(\rho_U)$ and c will be unmarked in $\eta' \bowtie \rho$ because of the choice of marking arcs in ρ_U. If instead $c \subseteq V$, then $c \subseteq U_{\eta' \bowtie \rho}$.

Since $U_{\mathrm{supp}(\nu)} = \varnothing$, this verifies $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\eta' \bowtie \rho)$ by Lemma 6.3.

Now assume ν has no unmarked base curves. Since ν is non-maximal, this means $U_{\mathrm{supp}(\nu)}$ must be non-empty. Since $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\eta) \cap \mathrm{lk}_B(\eta')$, we must have $U_{\mathrm{supp}(\nu)} \subseteq V_0$. If $U_{\mathrm{supp}(\nu)}$ contains an annular component of V_0, then every vertex of $\mathrm{lk}_B(\nu)$ would have to be equal to or joined by an edge to the core curve of that component. However, this cannot be the case since $\mathrm{lk}_B(\nu)$ is not a join. Thus, $U_{\mathrm{supp}(\nu)} \subseteq V = U_{\mathrm{supp}(\eta' \bowtie \rho)}$. Since ν has no unmarked base curves, Lemma 6.3 implies $\mathrm{lk}_B(\nu) \subseteq \mathrm{lk}_B(\eta' \bowtie \rho)$. □

Proposition 6.7. *Let μ be a non-maximal simplex of $B(X_\alpha)$ and let x, y be vertices of $\mathrm{lk}_B(\mu)$. If x and y are joined by a $B(W_\alpha)$-edge in $B(X_\alpha)^{+B(W_\alpha)}$ but not a $B(X_\alpha)$-edge, then there exists simplices ν_x and ν_y of $B(X_\alpha)$ so that $\mu \subseteq \nu_x, \nu_y$ and $\nu_x \cup x$ and $\nu_y \cup y$ are adjacent vertices of $B(W_\alpha)$.*

Proof. Assume first that x, y are joined by a $B(W_\alpha)$-edge coming from a twist edges of $B(W_\alpha)$. Then there exists $c \in X^0_\alpha$ so that $x, y \in B(c)$. Because $x, y \in \mathrm{lk}_B(\mu)$, the curve c must be in either $\mathrm{base}(\mu)$ or $\mathrm{lk}_X(\mathrm{base}(\mu))$. By Lemma 5.4, there exists a pants decomposition ρ of $S^z - c$ so that $\mathrm{base}(\mu) \cap (S^z - c) \subseteq \rho$ and $\rho \subseteq X_\alpha$. For each $d \in \rho$, select a marking arc $t_d \in B(d)$ so that t_d is the element of $\mathrm{markarc}(\mu) \cap B(d)$ whenever $\mathrm{markarc}(\mu) \cap B(d) \neq \varnothing$. Then, $\nu_x = \nu_y = \rho \cup \{t_d : d \in \rho\} \cup c$ are the desired simplices.

Now assume that x and y are joined by a $B(W_\alpha)$-edge coming from a flip edge of $B(W_\alpha)$. In this case, $\mathrm{supp}(x)$ and $\mathrm{supp}(y)$ fill a complexity 1 subsurface V of S and there exist marking arcs $t_x \in B(\mathrm{supp}(x))$ and $t_y \in B(\mathrm{supp}(y))$ so that $d_{A_{\mathrm{supp}(x)}}(t_x, \mathrm{supp}(y)) \leqslant 1$ and $d_{A_{\mathrm{supp}(y)}}(t_y, \mathrm{supp}(x)) \leqslant 1$. Such marking arc t_x, t_y are vertices of the maximal simplices in $B(X_\alpha)$ determined by the flip edge of $B(W_\alpha)$. Moreover x (resp. y) is equal to either t_x or $\mathrm{supp}(x)$ (resp. t_y or $\mathrm{supp}(y)$).

Since $x, y \in \mathrm{lk}_B(\mu)$, every curve of $\mathrm{base}(\mu)$ must be equal to or disjoint from each of $\mathrm{supp}(x)$ and $\mathrm{supp}(y)$. Hence $\mathrm{base}(\mu)$ does not intersect ∂V. There then exists a pants decomposition ρ of $S^z - V$ so that $\mathrm{base}(\mu) \cap (S^z - V) \subseteq \rho$ and $\rho \subseteq X_\alpha$ (Lemma 5.4). For each $c \in \rho \cup \partial V$, select a marking arc $t_c \in B(c)$ so that t_c is the element of $\mathrm{markarc}(\mu) \cap B(c)$ whenever $\mathrm{markarc}(\mu) \cap B(c) \neq \varnothing$. We now define a maximal simplex ν_x depending on the options for x.

- If $x = \text{supp}(x)$, then $\nu_x = \rho \cup \partial V \cup \{t_c : c \in \rho \cup \partial V\} \cup t_x$.
- If $x = t_x$, then $\nu_x = \rho \cup \partial V \cup \{t_c : c \in \rho \cup \partial V\} \cup \text{supp}(x)$.

We similarly define ν_y. Thus, $\nu_x \bowtie x$ and $\nu_y \bowtie y$ are joined by an edge of $B(W_\alpha)$. □

6.3 Reducing the geometric conditions to the non-annular case

Despite the addition of the marking arcs, the proof that $(B(X_\alpha), B(W_\alpha))$ satisfies the geometric conditions of a combinatorial HHS is nearly identical to the proof for (X_α, W_α). In this subsection, we will describe how to reduce the proof for $(B(X_\alpha), B(W_\alpha))$ to the case of (X_α, W_α).

To start, we prove the hyperbolicity of $B(X_\alpha)^{+B(W_\alpha)}$.

Lemma 6.8. *$B(X_\alpha)^{+B(W_\alpha)}$ is E_α-equivariantly quasi-isometric to both $B(X_\alpha)$ and X_α. Hence $B(X_\alpha)^{+B(W_\alpha)}$ and $B(X_\alpha)$ are both uniform quasi-trees.*

Proof. The definition of $B(W_\alpha)$ implies that the inclusion of $B(X_\alpha)$ into $B(X_\alpha)^{+B(W_\alpha)}$ is a quasi-isometry. This is because if two vertices x and y of $B(X_\alpha)^{+B(W_\alpha)}$ are joined by a $B(W_\alpha)$-edge and not a $B(X_\alpha)$-edge, then $x \in \mu$ and $y \in \nu$ where μ and ν are maximal simplices that are joined by an edge of $B(W_\alpha)$. However, μ and ν will have at least one base curve in common and this curve will need to be at most distance 1 from both x and y in $B(X_\alpha)$. Thus $d_{B(X_\alpha)}(x, y) \leqslant 2$.

Next we check that the inclusion of X_α into $B(X_\alpha)$ is also a quasi-isometry. Consider the map $B^0(X_\alpha) \to X_\alpha^0$ that sends $x \to \text{supp}(x)$. Lemma 6.1 showed this map is $(1, 0)$-coarsely Lipschitz, making the inclusion of X_α into $B(X_\alpha)$ an isometric embedding. Since each vertex of $B(X_\alpha)$ is at most 1 away from a vertex of X_α, the two spaces are quasi-isometric. □

We now prove the quasi-isometric embedding and hyperbolicity of the links of simplices of $B(X_\alpha)$. We will continue to use $\text{lk}_B(\mu)$ to denote the link of μ in $B(X_\alpha)$ and we will use $\text{BSat}(\mu)$ to denote the saturation of μ in $B(X_\alpha)$. Define $B(Y_\mu)$ to be $B(X_\alpha)^{+B(W_\alpha)} - \text{BSat}(\mu)$ and let $BH(\mu)$ be the subset of $B(Y_\mu)$ spanned by $\text{lk}_B(\mu)$—as before we are taking the link in $B(X_\alpha)$ then taking the span in $B(Y_\mu)$. The spaces $Y_{\text{base}(\mu)}$ and $H(\text{base}(\mu))$ are the spaces defined in Section 5.3. All graphs are given their intrinsic path metrics.

Proposition 6.9. *For each non-maximal simplex μ of $B(X_\alpha)$, the space $BH(\mu)$ is hyperbolic and the inclusion of $BH(\mu)$ into $B(Y_\mu)$ is a uniform quasi-isometric embedding. Moreover, $BH(\mu)$ falls into one of three cases:*

- *$BH(\mu)$ has bounded diameter because $\text{lk}_B(\mu)$ is a join or a single vertex.*
- *$\text{base}(\mu) = \text{supp}(\mu)$, $\text{base}(\mu)$ is not a pants decomposition of S^z, and μ contains no unmarked base curves. In this case, the inclusion of $H(\text{base}(\mu))$ into $BH(\mu)$ is a uniform quasi-isometry.*
- *$\text{base}(\mu)$ is a pants decomposition of S^z and μ contains exactly one unmarked base curve c. In this case, $BH(\mu)$ is equal to $\mathcal{C}(A_c)$.*

We break the proof into four cases based on μ. In the first two cases, the $B(X_\alpha)$-link of μ is join, so $BH(\mu)$ has bounded diameter, automatically satisfying the hyperbolicity and quasi-isometric embedding properties. In the third case, we reduce the problem to the results we established for $H(\text{base}(\mu))$ in the non-annular case. The final case addresses links consisting only of marking arcs. In this case, we use the subsurface projection to annular complexes.

Proof of Proposition 6.9.

1. Suppose $\text{supp}(\mu) \neq \text{base}(\mu)$. By Lemma 6.2, $\text{lk}_B(\mu)$ is either a join or a single vertex.

2. Suppose $\text{supp}(\mu) = \text{base}(\mu)$, $\text{base}(\mu)$ is not a pants decomposition of S^z, and there is an unmarked base curve $c \in \text{base}(\mu)$. In this case, every vertex of $\text{lk}_B(\mu) - B(c)$ is joined by an edge to each vertex of $B(c)$, so $\text{lk}_B(\mu)$ is a join.

3. Suppose $\text{supp}(\mu) = \text{base}(\mu)$, $\text{base}(\mu)$ is not a pants decomposition of S^z, and μ has no unmarked base curves. In this case, $\text{lk}_B(\mu)$ is a copy of $\text{lk}_X(\text{base}(\mu))$ with the $B(c)$-vertices attached to each X_α-vertex. Thus, we have an inclusion of $Y^0_{\text{base}(\mu)}$ into $B^0(Y_\mu)$ that sends $H^0(\text{base}(\mu))$ into $BH^0(\mu)$. Further, this map can be extended over edges as follows: if $x, y \in Y^0_{\text{base}(\mu)}$ are joined by an X_α-edge (i.e. they are disjoint curves), then they are joined by a $B(X_\alpha)$-edge as well. If instead x and y are joined by a W_α-edge, then the proof of Proposition 6.7 shows that we can also find two maximal simplices that are joined by a flip edge of $B(W_\alpha)$ and contain x and y respectively.

   We argue that this inclusion is actually a quasi-isometry on all of $Y_{\text{base}(\mu)}$ and when restricted to $H(\text{base}(\mu))$. Consider the map $B^0(Y_\mu) \to Y^0_{\text{base}(\mu)}$ defined by $y \to \text{supp}(y)$. An identical argument as given in Lemma 6.1 shows that this map induces a $(1, 0)$-coarsely Lipschitz map $B(Y_\mu) \to Y_{\text{base}(\mu)}$. Thus the inclusion $Y_{\text{base}(\mu)} \to B(Y_\mu)$ is an isometric embedding by Lemma 2.3. However, since each vertex of $B(Y_\mu)$ is at most 1 away from a vertex of $Y_{\text{base}(\mu)}$, the inclusion of $Y_{\text{base}(\mu)}$ into $B(Y_\mu)$ is a quasi-isometry. Further, the subset $H(\text{base}(\mu))$ is quasi-isometric to the subset $BH(\mu)$.

   Since $H(\text{base}(\mu))$ is hyperbolic and quasi-isometrically embeds in $Y_{\text{base}(\mu)}$ by Proposition 5.18, we have that $BH(\mu)$ is hyperbolic and quasi-isometrically embeds in $B(Y_\mu)$. Uniformity of the constants follows from the uniformity of the constants in Proposition 5.18.

4. Suppose $\text{supp}(\mu) = \text{base}(\mu)$ is a pants decomposition and $\text{unmark}(\mu) \neq \varnothing$. If μ has more than one unmarked base curve, then $\text{lk}_B(\mu)$ is a join of the $B(c)$ for these unmarked base curves and hence $BH(\mu)$ is bounded. Thus we can assume there exists exactly one unmarked base curve c. Hence, $\text{lk}_B(\mu) = B(c)$ and $BH(\mu)$ is the graph with vertex set $B(c)$ and an edge between to vertices t and t' if $d_{A_c}(t, t') = 1$. Since $B(c)$ is the vertex set of $\mathcal{C}(A_c)$, this means $BH(\mu) = \mathcal{C}(A_c)$ which is hyperbolic.

   We claim that every vertex $y \in B(Y_\mu) - BH(\mu)$ has $\text{supp}(y)$ intersecting c. First, we know $\text{supp}(y) \neq c$ as $\text{supp}(y) = c$ would imply that $y \in B(c) = BH^0(\mu)$. If $\text{supp}(y)$ was disjoint from c, then we could find a simplex ν of $B(X_\alpha)$ that contains y and has $\text{lk}_B(\nu) = B(c) = \text{lk}_B(\mu)$. To do this, start with a pants decomposition of $S^z - c$ that

contains $\mathrm{supp}(y)$ and uses only curves from X_α. Then select marking arcs for each of these curve so that y is a marking arc if $y \in B(\mathrm{supp}(y))$. Since the existence of such a simplex ν would imply $y \notin B(Y_\mu)$, it must be the case that $\mathrm{supp}(y)$ intersects c.

Since the curve $\mathrm{supp}(y)$ intersects c for each $y \in B(Y_\mu) - BH(\mu)$, the subsurface projection $\pi_{A_c}(\mathrm{supp}(y))$ is non-empty. Thus, we can define $\psi_\mu \colon B^0(Y_\mu) \to BH(\mu) = \mathcal{C}(A_c)$ by $\psi_\mu(y) = \pi_{A_c}(\mathrm{supp}(y))$ when $y \notin BH^0(\mu)$ and $\psi_\mu(y) = y$ when $y \in BH^0(\mu) = \mathcal{C}^0(A_c)$.

To prove ψ_μ is uniformly coarsely Lipschitz, we uniformly bound the distance between a pair of vertices of $B(Y_\mu)$ that are joined by an edge. If two curves $x, y \in B^0(Y_\mu) - BH^0(\mu)$ are joined by an edge of $B(Y_\mu)$ and $\mathrm{supp}(x) \neq \mathrm{supp}(y)$, then $i(\mathrm{supp}(x), \mathrm{supp}(y)) \leqslant 4$. By construction of the annular complex, $d_{A_c}(\mathrm{supp}(x), \mathrm{supp}(y))$ is uniformly bounded in terms of the intersection number $i(\mathrm{supp}(x), \mathrm{supp}(y))$. Thus $d_{BH(\mu)}(\psi_\mu(x), \psi(\mu))$ is uniformly bounded. If $y \in B^0(Y_\mu) - BH^0(\mu)$ and $x \in BH^0(\mu)$ are joined by an edge of $B(Y_\mu)$, then $d_{A_c}(\mathrm{supp}(y), x) \leqslant 1$ because they are joined by a $B(W_\alpha)$-edge coming from a flip move. The map ψ_μ is therefore uniformly coarsely Lipschitz, proving $BH(\mu)$ is quasi-isometrically embedded in $B(Y_\mu)$. □

6.4 E_α is an HHG

We now conclude our proof that E_α is a hierarchically hyperbolic group.

Theorem 6.10. *The pair $(B(X_\alpha), B(W_\alpha))$ is a combinatorial HHS and the group E_α acts on $B(X_\alpha)$ with finitely many orbits of links of simplices. In particular, E_α is a hierarchically hyperbolic group.*

Proof. We have proved the requirements for $(B(X_\alpha), B(W_\alpha))$ to be a combinatorial HHS in the following results:

- If $\mu_1, \ldots, \mu_n$ is a sequence of simplices of $B(X_\alpha)$ so that $\mathrm{lk}(\mu_1) \subsetneq \cdots \subsetneq \mathrm{lk}(\mu_n)$, then there is a corresponding sequence of properly nested subsurface of S^z by Lemma 6.3 and Remark 6.4. Hence, n is bounded in terms of the complexity of S^z, which is determined by S.
- The hyperbolicity of $B(X_\alpha)^{+B(W_\alpha)}$ (Item 2) was shown in Lemma 6.8.
- The hyperbolicity and quasi-isometric embedding conditions (Item 3) were verified in Proposition 6.9.
- The combinatorial conditions (Items 4 and 5) where shown in Propositions 6.5 and 6.7 respectively. Note, the condition proved is Proposition 6.5 is strictly stronger than what is required by the definition of a combinatorial HHS, since $\mathrm{diam}(BH(\mu)) > 2$ implies $\mathrm{lk}_B(\mu)$ is not a join.

This implies that $B(W_\alpha)$ is connected and a hierarchically hyperbolic space. We now verify the additional requirements from Theorem 2.18 for E_α to be a hierarchically hyperbolic group.

Because the action of E_α on $B(X_\alpha)$ is induced by the action (up to isotopy) of E_α on S^z, the action of E_α on the maximal simplices of $B(X_\alpha)$ induces the same action on $B(W_\alpha)$ as the action induced by E_α acting on S^z. It remains to verify that the action of E_α on $B(W_\alpha)$ is metrically proper and cobounded plus that there are finitely many E_α-orbits of links of simplices of $B(X_\alpha)$.

Coboundedly. Because E_α acts coboundedly on W_α it suffices to prove that for any $\mu, \nu \in B^0(W_\alpha)$, if $d_{W_\alpha}(\text{base}(\mu), \text{base}(\nu)) \leqslant 1$, then there exist $\phi, \psi \in E_\alpha$ so that

$$d_{B(W_\alpha)}(\phi(\mu), \psi(\nu))$$

is uniformly bounded.

Suppose $\mu, \nu \in B^0(W_\alpha)$ with $d_{W_\alpha}(\text{base}(\mu), \text{base}(\nu)) \leqslant 1$. Recall, there exists universal number $D \geqslant 1$ (independent of S) so that for each curve b on S^z, the quotient of the annular complex $\mathcal{C}(A_b)$ by the subgroup generated by the Dehn twist around b has diameter D. Thus, if $\text{base}(\mu) = \text{base}(\nu)$, then there exists $\phi \in \text{MCG}(S; z)$ so that ϕ is product of Dehn twists around curves in $\text{base}(\mu)$, and for each $b \in \text{base}(\mu)$, the marking arc of $\phi(\mu)$ for b is $\mathcal{C}(A_b)$-distance D from the marking arc of ν for b. Further, $\phi \in E_\alpha$ since it is the product of Dehn twist about curves in X_α. Hence, by performing at most D twist moves on each curve of ν, we can bound $d_{B(W_\alpha)}(\phi(\mu), \nu)$ by D times the number of curves in a pants decomposition of S^z as desired.

If instead $\text{base}(\mu)$ is joined by an edge of W_α to $\text{base}(\nu)$, then there must exist $c \in \text{base}(\mu)$ and $d \in \text{base}(\nu)$ so that $(\text{base}(\mu) - c) \cup d = \text{base}(\nu)$ and $i(c, d) \leqslant 4$. This implies $\pi_{A_c}(d) \subseteq \mathcal{C}^0(A_c) = B(c)$ and $\pi_{A_d}(c) \subseteq \mathcal{C}^0(A_d) = B(d)$.

Similarly to the previous case, there exists $\phi, \psi \in \text{MCG}(S; z)$ so that $\phi(\mu)$ and $\psi(\nu)$ satisfy:

1. $\text{base}(\phi(\mu)) = \text{base}(\mu)$ and $\text{base}(\psi(\nu)) = \text{base}(\nu)$ because ϕ and ψ are products of Dehn twists about curves in $\text{base}(\mu)$ and $\text{base}(\nu)$ respectively. Since each of these Dehn twists are around a curve in X_α, ϕ and ψ are both elements of E_α.

2. For each $b \in \text{base}(\mu) - c = \text{base}(\nu) - d$, the marking arc of $\phi(\mu)$ for b is $\mathcal{C}(A_b)$-distance D from the marking arc of $\psi(\nu)$ for b.

3. The marking arc of $\phi(\mu)$ for c is $\mathcal{C}(A_c)$-distance D from $\pi_{A_c}(d)$ and the marking arc of $\psi(\nu)$ for d is $\mathcal{C}(A_d)$-distance D from $\pi_{A_d}(c)$.

By performing at most D twist moves for each curve in $\text{base}(\phi(\mu))$ and at most D twist moves in the annular complex for $d \in \text{base}(\psi(\nu))$ we can produce two markings that differ by a flip between c and d. Thus, we have that $d_{B(W_\alpha)}(\phi(\mu), \psi(\nu))$ is bounded above by $2D$ times the number of curves in a pants decomposition of S^z.

Metrically properly. Our argument uses Masur and Minsky's graph of clean markings $\mathcal{M}(S^z)$ and the subsurface projection of markings. We direct the reader to Section 2.4 to recall the necessary definitions and properties that we will need for working with $\mathcal{M}(S^z)$.

For each $\mu \in B(W_\alpha)$, let $\text{cl}(\mu)$ be the set of clean markings that are compatible with μ; see Theorem 2.8. There is a uniform bound on the diameter of $\text{cl}(\mu)$ in $\mathcal{M}(S)$ depending on S^z. We first show that the map $B(W_\alpha) \to \mathcal{M}(S^z)$ by $\mu \to \text{cl}(\mu)$ is coarsely Lipschitz.

Let $\mu, \nu \in B(W_\alpha)$ be joined by an edge. The definition of the edges plus Lemma 2.4 ensures that $d_V(\mu, \nu)$ is uniformly bounded for every subsurface V of S^z. This implies $d_V(\mathrm{cl}(\mu), \mathrm{cl}(\nu))$ is uniformly bounded because for each $\mu' \in \mathrm{cl}(\mu)$ and $\nu' \in \mathrm{cl}(\nu)$ and for each subsurface V, we have

$$d_V(\mu', \nu') \leqslant d_V(\mu', \mu) + d_V(\mu, \nu) + d_V(\nu, \nu') + 6 \leqslant d_V(\mu, \nu) + 12.$$

By applying the second item in Theorem 2.8, we can therefore find $D \geqslant 0$, depending only on S^z, so that for any $\mu, \nu \in B(W_\alpha)$ joined by an edge we have

$$d_{\mathcal{M}(S)}(\mathrm{cl}(\mu), \mathrm{cl}(\nu)) \leqslant D.$$

Further, by possibly enlarging D, we can assume $\mathrm{diam}_{\mathcal{M}(S)}(\mathrm{cl}(\mu)) \leqslant D$ for each $\mu \in B(W_\alpha)$. By Lemma 2.2, the map $B(W_\alpha) \to \mathcal{M}(S^z)$ by $\mu \to \mathrm{cl}(\mu)$ is $(3D, 0)$-coarsely Lipschitz.

Now, let $\mathrm{ball}_r^B(\cdot)$ and $\mathrm{ball}_r^{\mathcal{M}}(\cdot)$ denote the balls of radius r in $B(W_\alpha)$ and $\mathcal{M}(S)$ respectively. For each $r \geqslant 0$ and $\mu \in B(W_\alpha)$, the set $\{g \in E_\alpha : g \cdot \mathrm{ball}_r^B(\mu) \cap \mathrm{ball}_r^B(\mu) \neq \varnothing\}$ is a subset of

$$\{g \in E_\alpha : g \cdot \mathrm{ball}_{3Dr}^{\mathcal{M}}(\mathrm{cl}(\mu)) \cap \mathrm{ball}_{3Dr}^{\mathcal{M}}(\mathrm{cl}(\mu)) \neq \varnothing\}.$$

Because E_α is a subgroup of $\mathrm{MCG}(S; z)$, E_α acts metrically properly on $\mathcal{M}(S)$ (Theorem 2.8). Thus,

$$\{g \in E_\alpha : g \cdot \mathrm{ball}_{3Dr}^{\mathcal{M}}(\mathrm{cl}(\mu)) \cap \mathrm{ball}_{3Dr}^{\mathcal{M}}(\mathrm{cl}(\mu)) \neq \varnothing\}$$

contains finitely many elements of E_α. This implies $\{g \in E_\alpha : g \cdot \mathrm{ball}_r^B(\mu) \cap \mathrm{ball}_r^B(\mu) \neq \varnothing\}$ also contains only finitely many elements of E_α, making the action of E_α on $B(W_\alpha)$ metrically proper.

Finitely many orbits of links. For a simplex μ of $B(X_\alpha)$, define $\overline{\mu} = \mathrm{supp}(\mu) \bowtie \mathrm{markarc}(\mu)$. Note, for any simplex μ, we have $\mathrm{base}(\overline{\mu}) = \mathrm{supp}(\overline{\mu}) = \mathrm{supp}(\mu)$. Because $\mathrm{base}(\overline{\mu}) = \mathrm{supp}(\overline{\mu})$, Lemma 6.3 implies that $\mathrm{lk}_B(\overline{\mu})$ is determined by the subsurface $U_{\mathrm{supp}(\overline{\mu})}$ and the unmarked base curves of μ. Because there are finitely many E_α-orbits of both curves in X_α and subsurfaces $V \subseteq S^z$ where $\partial V \subseteq X_\alpha$, there can only be a finite number of E_α-orbits of $B(X_\alpha)$-links of the $\overline{\mu}$'s. By Lemma 6.2, $\mathrm{lk}_B(\mu) - \mathrm{lk}_B(\overline{\mu}) = \mathrm{supp}(\mu) - \mathrm{base}(\mu)$ for each μ. Since $\mathrm{supp}(\mu) - \mathrm{base}(\mu)$ is a multicurve of curves in X_α, there are only finitely many E_α-orbits of possibilities for $\mathrm{supp}(\mu) - \mathrm{base}(\mu)$ as well. Thus, there are only finitely many E_α-orbits of $B(X_\alpha)$-links of simplices. □

7 A description of the HHG structure for E_α

We now provide an explicit description of the hierarchically hyperbolic group structure that E_α receives from Theorem 6.10. We start by recalling the defining information of a hierarchically hyperbolic structure (Sections 7.1). We embrace a slightly non-standard set of notation in this subsection to draw a distinction between curve complexes/subsurface projection maps and the hyperbolic spaces/projection maps of an abstract HHS. Next we describe the hierarchically hyperbolic structure that Theorem 2.18 imparts on an abstract combinatorial HHS (Section 7.2). The two subsequent subsections take this hierarchically hyperbolic structure for the space W_α (Section 7.3) and the group E_α (Section 7.4) and reinterpret it using the

topology of the surface . The final subsection uses this specific structure to show that some previously established results in the literature apply to E_α (Section 7.5).

This section is largely expository and intentionally does not provide the full definition of hierarchically hyperbolic spaces and groups. We direct the reader to [BHS19] and [Sis19] for detailed discussion of hierarchically hyperbolic spaces and to [BHMS20] for combinatorial hierarchically hyperbolic spaces.

7.1 The defining data of an HHS structure

An *E-hierarchically hyperbolic space structure* on a geodesic metric space $\mathcal{X}$ is a set $\mathfrak{S}$ indexing a collection of E-hyperbolic spaces $\{\mathcal{H}(V)\}_{V\in\mathfrak{S}}$ and equipped with three relations: nesting ($\sqsubseteq$), orthogonality ($\perp$), and transversality ($\pitchfork$). For each $V \in \mathfrak{S}$, there exists a surjective (E, E)-coarsely Lipschitz map $\tau_V\colon \mathcal{X} \to \mathcal{H}(V)$ called the *projection to* V. The relation $\sqsubseteq$ is a partial order on $\mathfrak{S}$ that contains a unique maximal element. Whenever $V \sqsubsetneq U$ in $\mathfrak{S}$ there exists a distinguished subset ρ^V_U in $\mathcal{H}(U)$ with diameter at most E. The relations $\perp$ and $\pitchfork$ are both symmetric and anti-reflexive. Whenever $V\pitchfork U$, there exist distinguished subsets $\rho^V_U \subseteq \mathcal{H}(U)$ and $\rho^U_V \subseteq \mathcal{H}(V)$ each of diameter at most E. When $V \sqsubsetneq U$ or $V\pitchfork U$, we call the subset ρ^V_U the *relative projection from* V *to* U. Every pair of distinct elements of $\mathfrak{S}$ is related by exactly one of the relations $\sqsubseteq$, $\perp$, or $\pitchfork$. To be a hierarchically hyperbolic space structure for $\mathcal{X}$, the set $\mathfrak{S}$ and these relations and projection need to satisfy a number of axioms; see [BHS19] or [Sis19] for a complete definition.

An HHS structure $\mathfrak{S}$ for a geodesic space $\mathcal{X}$ is an *hierarchically hyperbolic group structure* for a finitely generated group G if G acts metrically properly and coboundedly on $\mathcal{X}$ and there is an $\sqsubseteq$, $\perp$, and $\pitchfork$ preserving action of G on $\mathfrak{S}$ by bijections satisfying:

1. $\mathfrak{S}$ has finitely many G-orbits;
2. For each $V \in \mathfrak{S}$ and $g \in G$, there exists an isometry $g_V\colon \mathcal{H}(V) \to \mathcal{H}(gV)$ satisfying the following for all $V, U \in \mathfrak{S}$ and $g, h \in G$.
   - The map $(gh)_V\colon \mathcal{H}(V) \to \mathcal{H}(ghV)$ is equal to the map $g_{hV}\circ h_V\colon \mathcal{H}(V) \to \mathcal{H}(ghV)$.
   - For each $x \in \mathcal{X}$, $g_V(\tau_V(x))$ and $\tau_{gV}(g \cdot x)$ are at most E far apart in $\mathcal{H}(gV)$.
   - If $U\pitchfork V$ or $U \sqsubsetneq V$, then $g_V(\rho^U_V)$ and ρ^{gU}_{gV} are at most E far apart in $\mathcal{H}(gV)$.

7.2 The HHS structure for combinatorial HHSs

Let (X, W) be a combinatorial HHS. Recall the equivalence relation on simplices of X: $\Delta \sim \Delta'$ if $\mathrm{lk}(\Delta) = \mathrm{lk}(\Delta')$. Let $[\Delta]$ denote the $\sim$-equivalence class of the simplex Δ. We adopt the convention that the empty set is a simplex of X whose link is all of X. Let $\mathfrak{S}$ be the set of $\sim$-equivalence classes of all non-maximal simplices of X, including the empty set.

This set $\mathfrak{S}$ is the index set for the HHS structure for W provided by Theorem 2.18. We define $[\Delta] \sqsubseteq [\Delta']$ if $\mathrm{lk}(\Delta) \subseteq \mathrm{lk}(\Delta')$ and define $[\Delta] \perp [\Delta']$ if $\mathrm{lk}(\Delta) \subseteq \mathrm{lk}(\mathrm{lk}(\Delta'))$—or equivalently $\mathrm{lk}(\Delta') \subseteq \mathrm{lk}(\mathrm{lk}(\Delta))$. We define $[\Delta]\pitchfork[\Delta']$ if $[\Delta] \not\perp [\Delta']$ and neither is nested into the other. The $\sqsubseteq$-maximal element of $\mathfrak{S}$ is $[\varnothing]$ and its associated hyperbolic space is X^{+W}. For each $[\Delta] \in \mathfrak{S} - \{[\varnothing]\}$, the associated hyperbolic space is $H(\Delta)$. Note, this means the

hyperbolic space associated to $[\Delta]$ has bounded diameter whenever $\mathrm{lk}(\Delta)$ is a join or a single vertex.

The projection map $\tau_{[\Delta]}\colon W \to H(\Delta)$ is defined as follows. Let $\sigma \in W^0$. Since σ is a maximal simplex in X, there must be a vertex of σ contained in the space $Y_\Delta = X^{+W} - \mathrm{Sat}(\Delta)$. The space Y_Δ is shown to be hyperbolic in Section 3 of [BHMS20]. Since $H(\Delta)$ quasi-isometrically embeds into Y_Δ, there is a coarse closest point projection $\mathfrak{p}_\Delta\colon Y_\Delta \to H(\Delta)$. Define $\tau_{[\Delta]}(\sigma) = \mathfrak{p}_\Delta(\sigma \cap Y^0_\Delta)$. The relative projections are similarly defined by $\rho^{[\Delta']}_{[\Delta]} = \mathfrak{p}_\Delta(Y^0_{\Delta'} \cap Y^0_\Delta)$ whenever $[\Delta] \sqsubsetneq [\Delta']$ or $[\Delta] \pitchfork [\Delta']$. The hyperbolic spaces and projections are well defined because $Y_\Delta = Y_{\Delta'}$ and $H(\Delta) = H(\Delta')$ whenever $\Delta \sim \Delta'$.

7.3 The HHS structure for W_α

Let $\mathfrak{S}_\alpha$ be the HHS structure given to W_α by virtue of (X_α, W_α) being a combinatorial HHS. We give an interpretation of the relations and hyperbolic spaces for elements of $\mathfrak{S}_\alpha$ using the topology of subsurfaces and curves of S^z.

By Lemma 5.8, there is a bijection between links of simplices μ in X_α and the complimentary subsurfaces U_μ. Thus, we can identify the index set $\mathfrak{S}_\alpha$ with the set of subsurfaces $\{U_\mu : \mu \text{ is a non-maximal simplex of } X_\alpha\}$. Lemma 5.8 also says that the nesting of links $\mathrm{lk}(\mu) \subseteq \mathrm{lk}(\mu')$ is equivalent to the containment of subsurfaces $U_\mu \subseteq U_{\mu'}$. Lemma 7.1 below verifies that $\mathrm{lk}(\mu) \subseteq \mathrm{lk}(\mathrm{lk}(\mu'))$ if and only if U_μ and $U_{\mu'}$ are disjoint. Thus the nesting, orthogonality, and transversality of the elements of $\mathfrak{S}_\alpha$ respectively correspond to the containment, disjointness, and overlapping of the subsurfaces in $\{U_\mu : \mu \text{ is a non-maximal simplex of } X_\alpha\}$.

The hyperbolic spaces associated to subsurfaces come in two types. If μ is of subsurface type, then the hyperbolic space $H(\mu)$ is the modified curve graph $\mathcal{C}'(U_\mu)$ given in Definition 2.6. If μ is of tree guided type, then the hyperbolic space $H(\mu)$ is the subset of $\mathcal{C}'(S^z - \mu)$ spanned by the curves c with $\Pi(c)$ disjoint from α. This subset is quasi-isometric to tree.

Lemma 7.1. *For any non-maximal simplices $\mu, \mu' \subseteq X_\alpha$, we have $\mathrm{lk}(\mu) \subseteq \mathrm{lk}(\mathrm{lk}(\mu'))$ if and only if U_μ and $U_{\mu'}$ are disjoint*

Proof. Since $\mathrm{lk}(\mu)$ and $\mathrm{lk}(\mu')$ are spanned by the curves of X_α that are on U_μ and $U_{\mu'}$ respectively, if U_μ and $U_{\mu'}$ are disjoint, then $\mathrm{lk}(\mu) \subseteq \mathrm{lk}(\mathrm{lk}(\mu'))$ because every curve in $\mathrm{lk}(\mu)$ is disjoint from every curve in $\mathrm{lk}(\mu')$.

Conversely if $\mathrm{lk}(\mu) \subseteq \mathrm{lk}(\mathrm{lk}(\mu'))$, then U_μ must be disjoint from $U_{\mu'}$ since U_μ and $U_{\mu'}$ are filled by the curves in $\mathrm{lk}(\mu)$ and $\mathrm{lk}(\mu')$ respectively. □

7.4 The HHG structure for E_α

Let $\mathfrak{BS}_\alpha$ denote the HHS structure for $B(W_\alpha)$ arising from $(B(X_\alpha), B(W_\alpha))$ being a combinatorial HHS. By Theorem 6.10, $\mathfrak{BS}_\alpha$ is also an HHG structure for E_α. Let $\mathfrak{BS}^T_\alpha$ denote the set of equivalence classes of simplices $\mu \subseteq B(X_\alpha)$ with $\mathrm{supp}(\mu) = \mathrm{base}(\mu)$. We start by describing in topological terms the HHS data for equivalence classes of simplices in $\mathfrak{BS}^T_\alpha$. At the end we will describe why the other simplices can essentially be ignored when working with the HHG structure $\mathfrak{BS}_\alpha$ for E_α.

Define C_μ to be the disjoint union the set of annuli $\{A_c : c \in \mathrm{unmark}(\mu)\}$. Lemma 6.3 implies that $\mathrm{lk}_B(\mu) = \mathrm{lk}_B(\mu')$ if and only if $U_{\mathrm{supp}(\mu)} = U_{\mathrm{supp}(\mu')}$ and $C_\mu = C_{\mu'}$. Thus, the

set $\mathfrak{BS}_\alpha^T$ can be identified with the set of all pairs of subsurfaces $(U_{\mathrm{supp}(\mu)}, C_\mu)$. Lemma 6.3 showed that $\mathrm{lk}_B(\mu) \subseteq \mathrm{lk}_B(\mu')$, and hence the $\sqsubseteq$ relation, is equivalent to the following "dictionary nesting" of subsurfaces:

$$(U_{\mathrm{supp}(\mu)}, C_\mu) \sqsubseteq (U_{\mathrm{supp}(\mu')}, C_{\mu'}) \iff U_{\mathrm{supp}(\mu)} \subseteq U_{\mathrm{supp}(\mu')} \text{ and } C_\mu \subseteq U_{\mathrm{supp}(\mu')} \cup C_{\mu'}.$$

Lemma 7.2 below shows that $\mathrm{lk}_B(\mu) \subseteq \mathrm{lk}_B(\mathrm{lk}_B(\mu'))$ if and only if $U_{\mathrm{supp}(\mu)} \cup C_\mu$ is disjoint from $U_{\mathrm{supp}(\mu')} \cup C_{\mu'}$. Hence

$$(U_{\mathrm{supp}(\mu)}, C_\mu) \perp (U_{\mathrm{supp}(\mu')}, C_{\mu'}) \iff U_{\mathrm{supp}(\mu)} \cup C_\mu \text{ and } U_{\mathrm{supp}(\mu')} \cup C_{\mu'} \text{ are disjoint.}$$

Proposition 6.9 showed that there are three possibilities for the hyperbolic space $BH(\mu)$ when $[\mu] \in \mathfrak{BS}_\alpha$. When $U_{\mathrm{supp}(\mu)} \cup C_\mu$ is not connected, then $BH(\mu)$ has bounded diameter since $\mathrm{lk}_B(\mu)$ is a join. If $C_\mu = \varnothing$ and $U_{\mathrm{supp}(\mu)}$ is connected, then $BH(\mu)$ is quasi-isometric to $H(\mathrm{supp}(\mu))$, which in turn falls into one of the two case described in the previous section. Finally, when $U_{\mathrm{supp}(\mu)} = \varnothing$ and C_μ is an single annulus A_c, then $BH(\mu)$ is the annular complex $\mathcal{C}(A_c)$.

Finally consider a simplex μ where $\mathrm{base}(\mu) \neq \mathrm{supp}(\mu)$. For such a simplex, define $\overline{\mu} = \mathrm{supp}(\mu) \bowtie \mathrm{markarc}(\mu)$. Then $\mathrm{lk}_B(\mu) = \mathrm{lk}_B(\overline{\mu}) \bowtie (\mathrm{supp}(\mu) - \mathrm{base}(\mu))$. Thus, the hyperbolic spaces associated to equivalence classes of such simplices will have finite diameter. Further, for any simplex μ' if $\mathrm{lk}_B(\mu') \subseteq \mathrm{lk}_B(\mu)$, then either $\mathrm{lk}_B(\mu')$ is a join or $\mathrm{lk}_B(\mu') \subseteq \mathrm{lk}_B(\overline{\mu})$. With these two facts, one can essentially ignore the elements of $\mathfrak{BS}_\alpha$ with $\mathrm{base}(\mu) \neq \mathrm{supp}(\mu)$ when working with the hierarchically hyperbolic structure of $B(W_\alpha)$ (or E_α); Lemma 7.6 in the next section is an example of this philosophy. In fact, one could use these two facts to prove that $\mathfrak{BS}_\alpha^T$ itself an HHS structure for $B(W_\alpha)$. We have forgone this work as we have no need for it.

Lemma 7.2. *Let $\mu, \mu' \subseteq B(X_\alpha)$ be non-maximal simplices with $\mathrm{supp}(\mu) = \mathrm{base}(\mu)$ and $\mathrm{supp}(\mu') = \mathrm{base}(\mu')$. Then $\mathrm{lk}_B(\mu) \subseteq \mathrm{lk}_B(\mathrm{lk}_B(\mu'))$ if and only if $U_{\mathrm{supp}(\mu)} \cup C_\mu$ and $U_{\mathrm{supp}(\mu')} \cup C_{\mu'}$ are disjoint.*

Proof. We start by describing $\mathrm{lk}_B(\mathrm{lk}_B(\eta))$ when $\eta \subseteq B(X_\alpha)$ is a simplex with $\mathrm{supp}(\eta) = \mathrm{base}(\eta)$. By Lemma 6.2, $\mathrm{lk}_B(\eta)$ is spanned by the set of vertices:

$$\{x \in B(X_\alpha) : \mathrm{supp}(x) \in \mathrm{lk}_X(\mathrm{supp}(\eta))\} \cup \bigcup_{c \in \mathrm{unmark}(\eta)} B(c). \qquad (\star)$$

Let $y \in \mathrm{lk}_B(\mathrm{lk}_B(\eta))$. There are two facts we shall need about y.

I. $\mathrm{supp}(y)$ is disjoint from $U_{\mathrm{supp}(\eta)}$: to start, $\mathrm{supp}(y) \notin \mathrm{lk}_X(\mathrm{supp}(\eta))$, since $y \notin \mathrm{lk}_B(\eta)$. Thus, the only way for y to joined by edge to each vertex of $\mathrm{lk}_B(\eta)$, is for $\mathrm{supp}(y)$ to be disjoint from every curve in $\mathrm{lk}_X(\mathrm{supp}(\eta))$. Since $U_{\mathrm{supp}(\eta)}$ is defined to be the subsurface filled by the curves in $\mathrm{lk}_X(\mathrm{supp}(\eta))$, we must have that y is disjoint from $U_{\mathrm{supp}(\eta)}$.

II. for each unmarked base curve c of η, either $y = c$ or $\mathrm{supp}(y)$ is disjoint from c: if c is an unmarked base curve of η, then $B(c) \subseteq \mathrm{lk}_B(\eta)$. Hence, y must be joined by an edge to each vertex of $B(c)$. The only way for that to happen is if $y = c$ or $\mathrm{supp}(y)$ is disjoint from c.

Assume $\mathrm{lk}_B(\mu) \subseteq \mathrm{lk}_B(\mathrm{lk}_B(\mu'))$. Because of $(\star)$, $U_{\mathrm{supp}(\mu)} \cup C_\mu$ is precisely the subsurface filled by the set of curves $\{\mathrm{supp}(y) : y \in \mathrm{lk}_B(\mu)\}$. By Item I, $\mathrm{supp}(y)$ is disjoint from $U_{\mathrm{supp}(\mu')}$ for all $y \in \mathrm{lk}_B(\mu)$. Hence $U_{\mathrm{supp}(\mu)} \cup C_\mu$ is also disjoint from $U_{\mathrm{supp}(\mu')}$.

To show that $U_{\mathrm{supp}(\mu)} \cup C_\mu$ is disjoint from $C_{\mu'}$, we will prove that every curve in the set $\{\mathrm{supp}(y) : y \in \mathrm{lk}_B(\mu)\}$ is disjoint from every unmarked base curve of μ'. This implies $U_{\mathrm{supp}(\mu)} \cup C_\mu$ is disjoint from $C_{\mu'}$, since the unmarked base curves of μ' fill $C_{\mu'}$. By Item II, if $y \in \mathrm{lk}_B(\mu)$ and c' is an unmarked base curve of μ', then either $y = c'$ or $\mathrm{supp}(y)$ is disjoint from c'. Since $\mathrm{supp}(\mu) = \mathrm{base}(\mu)$, $(\star)$ implies that $B(c') \subseteq \mathrm{lk}_B(\mu)$ if $y = c'$. However, $B(c') \subseteq \mathrm{lk}_B(\mu')$ by $(\star)$, implying $B(c') \nsubseteq \mathrm{lk}_B(\mathrm{lk}_B(\mu'))$. Since $\mathrm{lk}_B(\mu) \subseteq \mathrm{lk}_B(\mathrm{lk}_B(\mu'))$, this means $y \neq c'$ and hence $\mathrm{supp}(y)$ must be disjoint from c' as desired.

Now assume that $U_{\mathrm{supp}(\mu)} \cup C_\mu$ and $U_{\mathrm{supp}(\mu')} \cup C_{\mu'}$ are disjoint and let $y \in \mathrm{lk}_B(\mu)$. Hence we have either $\mathrm{supp}(y) \subseteq U_{\mathrm{supp}(\mu)}$ or $\mathrm{supp}(y) \subseteq C_\mu$. Thus, $\mathrm{supp}(y)$ is disjoint from both $U_{\mathrm{supp}(\mu')}$ and $C_{\mu'}$. Since $\mathrm{supp}(y)$ is disjoint from $U_{\mathrm{supp}(\mu')}$, $\mathrm{supp}(y)$ is disjoint from every curve in $\mathrm{lk}_X(\mathrm{supp}(\mu'))$ and hence y joined be an edge to every vertex of

$$\{x \in B^0(X_\alpha) : \mathrm{supp}(x) \in \mathrm{lk}_X(\mathrm{supp}(\mu'))\}.$$

Similarly, $\mathrm{supp}(y)$ begin disjoint from $C_{\mu'}$ implies that $\mathrm{supp}(y)$ is disjoint from each unmarked base curve of μ'. But this implies y is joined by an edge to each vertex in $B(c')$ for each $c' \in \mathrm{unmark}(\mu')$. By $(\star)$, this makes $y \in \mathrm{lk}_B(\mathrm{lk}_B(\mu'))$ □

7.5 Some applications of the HHS structure

We now use the HHG structure $\mathfrak{BS}_\alpha$ to prove the remaining statements from Theorem 1.3 of the introduction. These results rely on two facts: (1) $X_\alpha^{+W_\alpha}$ is a quasi-tree (2) a minor modification of $\mathfrak{BS}_\alpha$ produces an HHG structure of E_α that has an additional property called *unbounded products* originally defined by Abbott, Behrstock, and Durham [ABD21]. The definition we give below of unbounded products is equivalent to the original definition in the setting of hierarchically hyperbolic groups and avoids having to describe unneeded additional aspects of the theory of hierarchically hyperbolic spaces. The proof of the equivalence is a straight forward application of the distance formula in an HHG [BHS19, Theorem 4.5] and the fact that G acts on $\mathfrak{S}$ with finitely many orbits.

Definition 7.3. A hierarchically hyperbolic group $(G, \mathfrak{S})$ has *unbounded products* if for all non-$\sqsubseteq$-maximal $U \in \mathfrak{S}$, whenever there exists $V \sqsubseteq U$ with $\mathrm{diam}(\mathcal{H}(V)) = \infty$, there also exists $Q \perp U$ so that $\mathrm{diam}(\mathcal{H}(Q)) = \infty$.

Given any hierarchically hyperbolic group $(G, \mathfrak{S})$, Abbott, Behrstock, and Durham provide an explicit construction of a new hierarchically hyperbolic group structure $\mathfrak{T}$ with unbounded products for G. In certain cases, this new structure maintains the same $\sqsubseteq$-maximal hyperbolic space.

Theorem 7.4 ([ABD21, Theorem 3.7])**.** *Let $(G, \mathfrak{S})$ be a hierarchically hyperbolic group. There exists a hierarchically hyperbolic groups structure $\mathfrak{T}$ for G with unbounded products. Moreover, if for every non-$\sqsubseteq$-maximal $U \in \mathfrak{S}$ with $\mathrm{diam}(\mathcal{H}(U)) = \infty$ there is $V \in \mathfrak{S}$ so that $V \perp U$ and $\mathrm{diam}(\mathcal{H}(V)) = \infty$, then the $\sqsubseteq$-maximal hyperbolic space of $\mathfrak{T}$ can be taken to be the $\sqsubseteq$-maximal hyperbolic space of $\mathfrak{S}$.*

Remark 7.5. While the moreover clause of Theorem 7.4 is not given in the statement in [ABD21], it is explicit in Abbott, Behrstock, and Durham's construction of the hyperbolic spaces for the new structure $\mathfrak{T}$.

Our HHG structure $\mathfrak{BS}_\alpha$ for E_α does not have unbounded products itself, but the next lemma verifies that $\mathfrak{BS}_\alpha$ does satisfy the hypotheses of the moreover clause of Theorem 7.4. This allows us to use Abbott, Behrstock, and Durham's construction without changing the $\sqsubseteq$-maximal hyperbolic space.

Lemma 7.6. *For every $[\mu] \in \mathfrak{BS}_\alpha$, if $\mu \neq \varnothing$ and $\operatorname{diam}(BH(\mu)) = \infty$, then there exist $[\nu] \in \mathfrak{S}_\alpha$ with $[\nu] \perp [\mu]$ and $\operatorname{diam}(BH(\nu)) = \infty$.*

Proof. Let μ be a non-maximal, non-empty simplex of $\mathfrak{BS}_\alpha$ so that $\operatorname{diam}(BH(\mu)) = \infty$. If $\operatorname{supp}(\mu) \neq \operatorname{base}(\mu)$, then Lemma 6.2 makes $\operatorname{lk}_B(\mu)$ a join. Since this would imply that $\operatorname{diam}(BH(\mu)) \leqslant 2$, we know $\operatorname{supp}(\mu) = \operatorname{base}(\mu)$.

Recall, C_μ is the disjoint union of annuli whose core curves are the unmarked base curves of μ. By Lemma 6.2, $U_{\operatorname{supp}(\mu)} \cup C_\mu$ is filled by the set of curves $\{\operatorname{supp}(y) : y \in \operatorname{lk}_B(\mu)\}$. Hence, if $U_{\operatorname{supp}(\mu)} \cup C_\mu$ is disconnected, then $\operatorname{lk}_B(\mu)$ would have to be a join. Since this would imply $BH(\mu)$ has bounded diameter, we know $U_{\operatorname{supp}(\mu)} \cup C_\mu$ is connected.

Since $U_{\operatorname{supp}(\mu)} \cup C_\mu$ is connected, we have two cases: either $C_\mu = \varnothing$ and $U_{\operatorname{supp}(\mu)}$ is connected or $U_{\operatorname{supp}(\mu)} = \varnothing$ and C_μ is an single annulus. In either case, there is a curve $c \in X^0_\alpha$ that is disjoint from $U_{\operatorname{supp}(\mu)} \cup C_\mu$ and is not the core curve of C_μ in the case $U_{\operatorname{supp}(\mu)} = \varnothing$. Lemma 5.4 implies that there exists a simplex σ of $B(X_\alpha)$ so that

- $c \in \operatorname{base}(\sigma)$;
- $\operatorname{base}(\sigma)$ is a pants decomposition of S^z;
- c is the only unmarked base curve of σ.

The $B(X_\alpha)$-link of σ is exactly $B(c)$, and hence $BH(\sigma)$ has infinite diameter as it is a quasi-isometric to a line (Proposition 6.9). Further, $[\sigma] \in \mathfrak{BS}^T_\alpha$, $U_{\operatorname{supp}(\sigma)} = \varnothing$, and C_σ is precisely the annulus with core curve c. Since C_σ is disjoint from $U_{\operatorname{supp}(\mu)} \cup C_\mu$ and $U_{\operatorname{supp}(\sigma)} = \varnothing$, we have $[\sigma] \perp [\mu]$ by Lemma 7.2. □

Abbott, Behrstock, and Durham prove the following results about HHGs with unbounded products; see [ABD21] for the relevant definitions.

Theorem 7.7 ([ABD21])**.** *Let $(G, \mathfrak{T})$ be a hierarchically hyperbolic group with unbounded products and let T be the $\sqsubseteq$-maximal element of $\mathfrak{T}$.*

- *The action of G on $\mathcal{H}(T)$ is the largest cobounded acylindrical action of G on a hyperbolic space.*
- *A subgroup $H < G$ is stable if and only if the orbit map of H into $\mathcal{H}(T)$ is a quasi-isometric embedding.*

Cordes, Charney, and Sisto recently provided a complete topological characterization of the Morse boundaries of certain groups [CCS23]. While they do not tackle hierarchically hyperbolic groups directly, their methods are readily adapted to certain hierarchically hyperbolic groups with unbounded products. In the appendix, we use their techniques along with the work of Abbott, Behrstock, and Durham to establish the following; see the appendix or [CCS23] for the definition of the Morse boundary and ω-Cantor space.

Theorem 7.8. *Let $(G, \mathfrak{T})$ be a hierarchically hyperbolic group with unbounded products and let T be the $\sqsubseteq$-maximal element of $\mathfrak{T}$. If $\mathcal{H}(T)$ is a quasi-tree and G is not hyperbolic, then the Morse boundary of G is homeomorphic to an ω-Cantor space.*

Proof. This is Corollary A.8 in the appendix. □

Combining these results, we can provide a proof of the remaining claims of Theorem 1.3.

Corollary 7.9.

1. *The action of E_α on X_α is the largest cobounded acylindrical action of E_α on a hyperbolic space.*

2. *A subgroup $H < E_\alpha$ is stable if and only if the orbit map of H into X_α is a quasi-isometric embedding. In particular, every stable subgroup of E_α is virtually free.*

3. *The Morse boundary of E_α is an ω-Cantor set.*

Proof. Combining Lemma 7.6 with Theorem 7.4, there is an HHG structure $\mathfrak{T}$ for E_α where the $\sqsubseteq$-maximal hyperbolic space of $\mathfrak{T}$ is the $\sqsubseteq$-maximal hyperbolic space of $\mathfrak{BS}_\alpha$. Lemma 6.8 showed that the $\sqsubseteq$-maximal element of $\mathfrak{BS}_\alpha$ is E_α-equivariantly quasi-isometric to the quasi-tree X_α. The conclusions then follow from Theorems 7.7 and 7.8 plus the fact that any finitely generated group whose orbit map in a quasi-tree is a quasi-isometric embedding is virtually free. □

A Appendix: HHGs and ω-Cantor set boundaries

Charney, Cordes, and Sisto recently gave the first topological descriptions of the Morse boundaries of non-hyperbolic groups [CCS23]. They proved the Morse boundaries of most right-angled Artin groups and graph manifold groups are homeomorphic to a specific limit of Cantor spaces that they call an ω-Cantor space. This appendix describes how their techniques can be straight forwardly extended to prove the same result for a broader class of groups. The main motivation is to show that certain hierarchically hyperbolic groups will also have ω-Cantor space boundaries. However, as the full power of hierarchical hyperbolicity will not be needed for our proof, we will instead work in the simpler setting of *Morse detectable* groups introduced by the author, Spriano, and Tran [RST22].

Notation

Throughout, G will be a finitely generated group with a fixed finite generating set. Whenever we discuss G as a metric space, we are referring to the path metric on the Cayley graph of G with respect to this finite generating set. We will denote the Morse boundary of G by $\partial_M G$ and the Gromov boundary of a hyperbolic space X by ∂X. Throughout, I will denote a closed interval of the real line. We remark that we are intentionally glib about the theory of Morse boundaries and direct the reader to Cordes' excellent survey on the topic for more details [Cor19].

Morse Boundaries

Given a function $N\colon [1,\infty)\times[0,\infty)\to[0,\infty)$ we say that a geodesic $\gamma\colon I\to G$ is N-Morse if every (k,c)-quasi-geodesic α with endpoints on γ is contained in the $N(k,c)$-neighborhood of γ. We call the function N the *Morse gauge* of γ. Let $\mathcal{M}$ denote the set of Morse gauges of all geodesics in G. The set $\mathcal{M}$ has a partial order where $N\preccurlyeq N'$ if $N(k,c)\leqslant N'(k,c)$ for all $k\geqslant 1$ and $c\geqslant 0$.

For each $N\in\mathcal{M}$, we let G^N denote the set of elements $g\in G$ so that there exists a N-Morse geodesic connecting the identity e to g. Cordes and Hume proved that the restriction of the metric of G to G^N makes G^N into a δ-hyperbolic metric space, where δ is determined by N [CD17, Proposition 2.3]. Let $\partial_M^N G$ denote the Gromov boundary of G^N with e as the base point. By definition, if $N\preccurlyeq N'$, then $G^N\subseteq G^{N'}$ and $\partial_M^N G\subseteq\partial_M^{N'}G$.

The *Morse boundary* of G is the direct limit of the $\partial_M^N G$. That is,

$$\partial_M G=\varinjlim_{\mathcal{M}}\partial_M^N G.$$

A consequence of this definition is that each strata $\partial_M^N G$ is a compact subset of the Morse boundary.

Lemma A.1 ([CH17, Theorem 3.14])**.** *For each Morse gauge N, $\partial_M^N G$ is a compact subset of $\partial_M G$.*

An important property of the Morse boundary is that the Gromov boundaries of stable subgroup embed into the Morse boundary of the entire group; see [CD17] or [DT15] for the definition of a stable subgroup.

Lemma A.2 ([CD17, Corollary 2.12])**.** *If H is a stable subgroup of G, then H is hyperbolic and the inclusion of H into G induces a continuous injection of ∂H into $\partial_M G$.*

Charney, Cordes, and Sisto defined an ω-Cantor space as the direct limit of countably many Cantor spaces where each Cantor space has empty interior in the next Cantor space; see [CCS23] for full details. They prove that all ω-Cantor sets are homeomorphic and give the following sufficient condition for the Morse boundary of a group to be an ω-Cantor set.

Theorem A.3 ([CCS23, Theorem 1.4])**.** *If $\partial_M G$ is totally disconnected, σ-compact, and contains a Cantor space, then $\partial_M G$ is homeomorphic to either a Cantor space or an ω-Cantor space. Moreover, $\partial_M G$ is a Cantor space if and only if G is virtually free.*

Morse detectability

Our proof that certain HHGs have ω-Cantor space boundaries only needs the following *Morse detectable* property.

Definition A.4 (Morse Detectable)**.** A finitely generated group G is *Morse detectable* if there exists a hyperbolic space X and a coarsely Lipschitz map $\pi\colon G \to X$ so that for any (k, c)-quasi-geodesic $\gamma\colon I \to G$ the following hold.

1. For each $N \in \mathcal{M}$, if γ is N-Morse, then there exists $\lambda \geqslant 1$ depending on N, k, and c, so that $\pi \circ \gamma$ is a (λ, λ)-quasi-geodesic of X.

2. For each $\lambda \geqslant 1$, if $\pi \circ \gamma$ is an (λ, λ)-quasi-geodesic, then γ is N-Morse for some N determined by k, c, and λ.

If the space X is a quasi-tree, then we say G is *Morse detectable in a quasi-tree*.

Abbott, Behrstock, and Durham prove that unbounded products imply a hierarchically hyperbolic groups is also Morse detectable.

Theorem A.5 ([ABD21, Corollary 6.2])**.** *If $(G, \mathfrak{S})$ is a hierarchically hyperbolic group with unbounded products, then G is Morse detectable in the $\sqsubseteq$-maximal hyperbolic space of $\mathfrak{S}$.*

Morse detectability allows us to conclude that the Morse boundary is σ-compact and continuously embeds in the Gromov boundary of the detecting space.

Lemma A.6. *If G is Morse detectable in the space X, then*

1. *$\partial_M G$ is σ-compact;*

2. *the map $\pi\colon G \to X$ induces a continuous injection $\partial\pi\colon \partial_M G \to \partial X$.*

Proof. **σ-compact.** For each $\lambda \in \mathbb{N}$, define $\partial_M^\lambda G$ to be the subset of points in $\partial_M G$ that are represented by a geodesic ray $\gamma\colon I \to G$ so that $\pi \circ \gamma$ is a (λ, λ)-quasi-geodesic ray in X. We claim that the closures of the $\partial_M^\lambda G$ give an exhaustion of $\partial_M G$ by compact sets.

Because G is Morse detectable, for each $\lambda \in \mathbb{N}$, there exists $N \in \mathcal{M}$ so that $\partial_M^\lambda G$ is contained in some $\partial_M^N G$ for N determined by λ. Thus the closure of each $\partial_M^\lambda G$ is compact since each $\partial_M^N G$ is compact (Lemma A.1). On the other hand, Morse detectability also implies that for each Morse gauge $N \in \mathcal{M}$, there is some $\lambda \in \mathbb{N}$ so that $\partial_M^N G \subseteq \partial_M^\lambda G$. Thus, $\partial_M G$ is contained in the union of the closures of the $\partial_M^\lambda G$.

Continuous embedding. The definition of Morse detectable make the restriction of π to G^N a quasi-isometric embedding with constants depending ultimately only on N. Since quasi-isometric embeddings of hyperbolic spaces extend to continuous injections on their boundaries, we have that π induces a continuous injection $\partial_M^N G \to \partial X$ for each $N \in \mathcal{M}$. Since $\partial_M G$ is the direct limit of the $\partial_M^N G$'s, this induces a continuous injection $\partial\pi\colon \partial_M G \to \partial X$. □

When our group is Morse detectable in a quasi-tree, we can verify Charney, Cordes, and Sisto's condition for the boundary to be an ω-Cantor set.

Theorem A.7. *Suppose $\partial_M G$ is non-empty. If G is Morse detectable in a quasi-tree and not hyperbolic, then $\partial_M G$ is homeomorphic to an ω-Cantor space.*

Proof. Let X be the quasi-tree that detects the Morse quasi-geodesics of G. By Lemma A.6, $\partial_M G$ is σ-compact. Since the Gromov boundary of a quasi-tree is totally disconnected, the continuous injection $\partial\pi\colon \partial_M G \to \partial X$ forces $\partial_M G$ to be totally disconnected. Therefore, the only requirement of Theorem A.3 that we still need to verify is that $\partial_M G$ contains a Cantor space.

Corollary 4.9 of [RST22] says any Morse detectable group that has non-empty Morse boundary is either virtually cyclic or contains a stable free subgroup of rank at least 2. The desired Cantor subspace thus comes from the Gromov boundary of this stable free subgroup by Lemma A.2. □

As a corollary of Theorem A.7, we show that the Morse boundaries of some HHGs are ω-Cantor spaces. This proves Theorem 7.8 from the main text.

Corollary A.8. *Let $(G, \mathfrak{S})$ be a non-hyperbolic hierarchically hyperbolic group with non-empty Morse boundary. If $\mathfrak{S}$ has unbounded products and the $\sqsubseteq$-maximal hyperbolic space of $\mathfrak{S}$ is a quasi-tree, then $\partial_M G$ is homeomorphic to an ω-Cantor space.*

Proof. By Theorem A.5, the group G is Morse detectable with the $\sqsubseteq$-maximal hyperbolic space in $\mathfrak{S}$. Thus, the conclusion is a special case of Theorem A.7. □

We conclude by highlighting another hierarchically hyperbolic group whose Morse boundary is an ω-Cantor space because of Corollary A.8.

Example A.9. Let H be the handlebody group of a genus 2 handlebody. Che sser proved H has a hierarchically hyperbolic group structure with unbounded products [Che22]. Further, Chesser shows that the $\sqsubseteq$-maximal space in this structure is a quasi-tree because it is the contact graph of a CAT(0) cube complex (they also prove this space is quasi-isometric to the disk graph of the genus 2 handlebody). Thus, $\partial_M H$ is an ω-Cantor space. In the case of genus 3 or larger, the handlebody group is known to not be hierarchically hyperbolic [HH21] although the question of whether it is or is not Morse detectable is open.

References

[ABD21] Carolyn Abbott, Jason Behrstock, and Matthew G. Durham. Largest acylindrical actions and stability in hierarchically hyperbolic groups. With an appendix by Daniel Berlyne and Jacob Russell. *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.

[Aou13] Tarik Aougab. Uniform hyperbolicity of the graphs of curves. *Geom. Topol.*, 17(5):2855–2875, 2013.

[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.

[BDM09] Jason Behrstock, Cornelia Druţu, and Lee Mosher. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. *Math. Ann.*, 344(3):543–595, 2009.

[Beh06] Jason A. Behrstock. Asymptotic geometry of the mapping class group and Teichmüller space. *Geom. Topol.*, 10:1523–1578, 2006.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BHMS20] Jason Behrstock, Mark Hagen, Alexandre Martin, and Alessandro Sisto. A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups. arXiv: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 F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: combination theorems and the distance formula. *Pacific J. Math.*, 299:257–338, 2019.

[BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.*, 170:909–996, 2021.

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

[BM08] Jason A. Behrstock and Yair N. Minsky. Dimension and rank for mapping class groups. *Ann. of Math. (2)*, 167(3):1055–1077, 2008.

[Bow06] Brian H. Bowditch. Intersection numbers and the hyperbolicity of the curve complex. *J. Reine Angew. Math.*, 598:105–129, 2006.

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

[BR20] Federico Berlai and Bruno Robbio. A refined combination theorem for hierarchically hyperbolic groups. *Groups Geom. Dyn.*, 14(4):1127–1203, 2020.

[Bro03] Jeffrey F. Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *J. Amer. Math. Soc.*, 16(3):495–535, 2003.

[CCS23] Ruth Charney, Matthew Cordes, and Alessandro Sisto. Complete topological descriptions of certain Morse boundaries. *Groups Geom. Dyn.*, 17(1):157–184, 2023.

[CD17] Matthew Cordes and Matthew G. Durham. Boundary convex cocompactness and stability of subgroups of finitely generated groups. *International Mathematics Research Notices*, 2019(6):1699–1724, 2017.

[CH17] Matthew Cordes and David Hume. Stability and the Morse boundary. *J. Lond. Math. Soc. (2)*, 95(3):963–988, 2017.

[Che22] Marissa Chesser. Stable subgroups of the genus 2 handlebody group. *Algebr. Geom. Topol.*, 22(2):919–971, 2022.

[Cor19] Matthew Cordes. A survey on Morse boundaries and stability. In *Beyond hyperbolicity*, volume 454 of *London Math. Soc. Lecture Note Ser.*, pages 83–116. Cambridge Univ. Press, Cambridge, 2019.

[CRS14] Matt Clay, Kasra Rafi, and Saul Schleimer. Uniform hyperbolicity of the curve graph via surgery sequences. *Algebr. Geom. Topol.*, 14(6):3325–3344, 2014.

[DDLS] Spencer Dowdal, Matthew G. Durham, Christopher J. Leininger, and Alessandro Sisto. Extensions of Veech groups II: Hierarchical hyperbolicity and quasi-isometric rigidity. arXiv:2111.00685.

[DT15] Matthew G. Durham and Samuel J. Taylor. Convex cocompactness and stability in mapping class groups. *Algebr. Geom. Topol.*, 15(5):2839–2859, 2015.

[Dur16] Matthew G. Durham. The augmented marking complex of a surface. *Journal of the London Mathematical Society*, 94(3):933, 2016.

[EMR17] Alex Eskin, Howard A. Masur, and Kasra Rafi. Large-scale rank of Teichmüller space. *Duke Math. J.*, 166(8):1517–1572, 2017.

[FM02] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152, 2002.

[Ham] Ursula Hamenstadt. Word hyperbolic extensions of surface groups. arXiv:math/0505244, 2005.

[HH21] Ursula Hamenstädt and Sebastian Hensel. The geometry of the handlebody groups II: Dehn functions. *Michigan Math. J.*, 70(1):–, 2021.

[HHP20] Thomas Haettel, Nima Hoda, and Harry Petyt. The coarse helly property, hierarchical hyperbolicity, and semihyperbolicity. arXiv:2009.14053, 2020.

[HMS] Mark Hagen, Alexandre Martin, and Alessandro Sisto. Artin groups of extra large type are hierarchically hyperbolic. arXiv:2109.04387, 2021.

[HPW15] Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb. 1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs. *J. Eur. Math. Soc. (JEMS)*, 17(4):755–762, 2015.

[HS20] Mark F. Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. *Israel J. Math.*, 236(1):45–89, 2020.

[KL07] Autumn 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.

[KLS09] Autumn E. Kent, Christopher J. Leininger, and Saul Schleimer. Trees and mapping class groups. *J. Reine Angew. Math.*, 637:1–21, 2009.

[KR14] Ilya Kapovich and Kasra Rafi. On hyperbolicity of free splitting and free factor complexes. *Groups Geom. Dyn.*, 8(2):391–414, 2014.

[Man05] Jason F. Manning. Geometry of pseudocharacters. *Geom. Topol.*, 9(2):1147–1185, 2005.

[MM99] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999.

[MM00] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.

[MS12] Mahan Mj and Pranab Sardar. A combination theorem for metric bundles. *Geom. Funct. Anal.*, 22(6):1636–1707, 2012.

[Ota96] Jean-Pierre Otal. Le théorème d'hyperbolisation pour les variétés fibrées de dimension 3. *Astérisque*, (235):x+159, 1996.

[Raf07] Kasra Rafi. A combinatorial model for the Teichmüller metric. *Geom. Funct. Anal.*, 17(3):936–959, 2007.

[RS] Bruno Robbio and Davide Spriano. Hierarchical hyperbolicity of hyperbolic-2-decomposable groups. arXiv:2007.13383.

[RST22] Jacob Russell, Davide Spriano, and Hung Cong Tran. The local-to-global property for Morse quasi-geodesics. *Math. Z.*, 300(2):1557–1602, 2022.

[RST23] Jacob Russell, Davide Spriano, and Hung C. Tran. Convexity in hierarchically hyperbolic spaces. *Algebr. Geom. Topol.*, 23(3):1167–1248, 2023.

[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.

[Vok22] Kate M. Vokes. Hierarchical hyperbolicity of graphs of multicurves. *Algebr. Geom. Topol.*, 22(1):113–151, 2022.