

PSEUDO-ANOSOV SUBGROUPS OF GENERAL FIBERED 3-MANIFOLD GROUPS

CHRISTOPHER J. LEININGER AND JACOB RUSSELL

ABSTRACT. We show that finitely generated and purely pseudo-Anosov subgroups of fundamental groups of fibered 3-manifolds with reducible monodromy are convex cocompact as subgroups of the mapping class group via the Birman exact sequence. Combined with results of Dowdall–Kent–Leininger and Kent–Leininger–Schleimer, this establishes the result for the image of all such fibered 3-manifold groups in the mapping class group.

1. INTRODUCTION

Farb and Mosher defined convex cocompactness in $\text{Mod}(S)$ —the mapping class group of a finite type orientable surface S of negative Euler characteristic—via analogy with convex cocompactness of Kleinian groups [FM02]. The convex cocompact subgroups of $\text{Mod}(S)$ play an important role in the geometry of surface group extensions and surfaces bundles [FM02, Ham, MS12] and have a rich dynamical and geometric structure [Ham07, KL08a, KL08b, DT15, BBKL20]. One basic property is that convex cocompact subgroups of $\text{Mod}(S)$ are finitely generated and *purely pseudo-Anosov*, that is, every infinite order element is pseudo-Anosov. In their introductory paper, Farb and Mosher asked if this pair of properties characterized convex cocompactness [FM02, Question 1.5]. A “no” answer can be used to produce a relatively simple example of a finitely generated group that is not hyperbolic, but has no Baumslag–Solitar subgroups, see [KL07, §8]. While, a “yes” answer would limit the possibilities for convex compact subgroups by requiring every finitely generated subgroup of such a group to again be convex cocompact. We establish that the answer to Farb and Mosher’s question is yes for subgroups that are contained in the image of the fundamental groups of fibered 3-manifold groups inside the mapping class group, as we now explain.

Every orientable 3-manifold that fibers over a circle is the mapping torus $M_f = S \times [0, 1] / (x, 1) \sim (f(x), 0)$ of an orientation preserving surface homeomorphism $f: S \rightarrow S$. Fixing a basepoint $z \in S \subset M_f$, the fundamental group, $\Gamma_f = \pi_1 M_f = \pi_1(M_f, z)$, splits as a semi-direct product, $\Gamma_f \cong \pi_1 S \rtimes_{f_*} \mathbb{Z}$, where f_* is an automorphism of $\pi_1 S = \pi_1(S, z)$ induced by f . If $\phi: \Gamma_f \rightarrow \mathbb{Z}$ is the homomorphism of this splitting, then we write $\mu: \mathbb{Z} \rightarrow \langle f \rangle < \text{Mod}(S)$ for the monodromy homomorphism, so that $\mu(\phi(g)) = f^{\phi(g)}$ for $g \in \Gamma_f$. Setting $S^z = S \setminus \{z\}$, the monodromy is the descent of a homomorphism $\mu^z: \Gamma_f \rightarrow \text{Mod}(S^z)$, with image in the finite index subgroup $\text{Mod}(S^z, z) < \text{Mod}(S^z)$ consisting of isotopy classes of homeomorphisms that fix the z -puncture. These homomorphisms fit into a commutative diagram with the *Birman exact sequence*, defined by the homomorphism $\Phi_*: \text{Mod}(S^z, z) \rightarrow \text{Mod}(S)$ induced by the inclusion $\Phi: S^z \rightarrow S$:

$$\begin{array}{ccccccc} 1 & \longrightarrow & \pi_1 S & \longrightarrow & \Gamma_f & \xrightarrow{\phi} & \mathbb{Z} \longrightarrow 1 \\ & & \downarrow = & & \downarrow \mu^z & & \downarrow \mu \\ 1 & \longrightarrow & \pi_1 S & \longrightarrow & \text{Mod}(S^z, z) & \xrightarrow{\Phi_*} & \text{Mod}(S) \longrightarrow 1. \end{array}$$

The Nielsen–Thurston Classification Theorem [FLP91] says that every element of a mapping class group is either pseudo-Anosov, reducible, or finite order. Our main result proves that the answer to Farb and Mosher’s question is “yes” for subgroups of $\mu^z(\Gamma_f)$ when f is an infinite order, reducible mapping class.

Theorem 1.1. *Suppose $\chi(S) < 0$ and $f: S \rightarrow S$ is a reducible, infinite order mapping class. For any subgroup $G < \Gamma_f$, the group $\mu^z(G) < \text{Mod}(S^z)$ is convex cocompact if and only if it is finitely generated and purely pseudo-Anosov.*

Remark 1.2. We note that any finitely generated, purely pseudo-Anosov subgroup $G < \Gamma_f$ as in Theorem 1.1 is necessarily free; see Lemma 3.3.

The analogue of Theorem 1.1 when f is pseudo-Anosov was previously shown to be true in [DKL14]. The analogue for f finite order is a consequence of the result for subgroups $G < \pi_1 S$, proved in [KLS09, Theorem 6.1]. Indeed, in this case $\mu^z(\Gamma_f)$ contains $\pi_1 S$ with finite index, and convex cocompactness is preserved by passage to finite index super- and subgroups. Combining these results, we see that the conclusion holds for any $f \in \text{Mod}(S)$.

Theorem 1.3. *Suppose $\chi(S) < 0$ and let $f: S \rightarrow S$ be any mapping class. For any subgroup $G < \Gamma_f$, the group $\mu^z(G) < \text{Mod}(S^z)$ is convex cocompact if and only if it is finitely generated and purely pseudo-Anosov.*

1.1. Known results. There are a number of other settings where finitely generated, purely pseudo-Anosov subgroups have been shown to be convex cocompact, providing an affirmative answer to Farb and Mosher’s question [FM02, Question 1.5]. As mentioned above, if G is a subgroup of either $\pi_1 S$ or Γ_f for f pseudo-Anosov, then the inclusion of G into $\text{Mod}(S^z)$ via the Birman exact sequence is convex cocompact if and only if it is finitely generated and purely pseudo-Anosov [KLS09, DKL14]. The same result has been proved under the assumption that G is either a subgroup of an *admissibly embedded*¹ right-angled Artin subgroup $A < \text{Mod}(S)$ [KMT17] or if G is contained in the genus-2 Goeritz group [Tsh21]. In [DT15], it was shown that $G < \text{Mod}(S)$ is convex cocompact if and only if G is a stable subgroup (except for two sporadic surfaces S), providing a purely geometric group theoretic characterization. This was strengthened in [BBKL20] where it was shown that $G < \text{Mod}(S)$ is convex cocompact if and only if G is finitely generated, purely pseudo-Anosov, and undistorted.

1.2. Proof summary. When $f: S \rightarrow S$ has infinite order, $\mu^z: \Gamma_f \rightarrow \text{Mod}(S^z, z)$ is injective, and we identify Γ_f with $\mu^z(\Gamma_f)$. For simplicity, we assume S is closed in this summary, ensuring that $\Phi: S^z \rightarrow S$ sends every essential curve on S^z to an essential curve on S . To prove Theorem 1.1, we fix a finitely generated and purely pseudo-Anosov $G < \Gamma_f$ and show that the orbit map of G to the curve complex $\mathcal{C}(S^z)$ is a quasi-isometric embedding. From [Ham07, KL08a], this is equivalent to G being convex cocompact; see Theorem 2.4. The central task is then to find a way to relate distances in G to distances in the curve complex.

For subgroups $G < \pi_1 S$, such a relationship was established in [KLS09] by examining K_u , the stabilizer in $\pi_1 S < \text{Mod}(S^z)$ of a simplex $u \subset \mathcal{C}(S^z)$. Using the isometric action of $\pi_1 S$ by deck transformations on the universal cover $p: \mathbb{H}^2 \rightarrow S$, we define $\mathfrak{H}_u$ to be the convex hull of the limit set of K_u in $\partial\mathbb{H}^2$. The group G also has a convex hull for its limit set, $\mathfrak{H}_G$, on which it will act geometrically and which therefore serves as a geometric model for G . A key argument in [KLS09] proves that $\mathfrak{H}_G \cap \mathfrak{H}_u$ has uniformly bounded diameter, independent u . The simplices that make up a geodesic edge path between G -orbit points in $\mathcal{C}(S^z)$ then give rise to a chain of bounded diameter sets in $\mathfrak{H}_G$. The total diameter of this chain bounds distance by a linear function of distance in $\mathcal{C}(S^z)$, as required. A similar approach is used in [DKL14] for $G < \Gamma_f$, when f is a pseudo-Anosov element of $\text{Mod}(S)$. In this case, the mapping torus M_f is a hyperbolic 3-manifold, thus the convex hulls for G and for simplex stabilizers can be taken in $\mathbb{H}^3$ instead of $\mathbb{H}^2$. Once again, the key result is that these convex hulls intersect in uniformly bounded diameter sets.

Our proof in the reducible case is inspired by these methods. The first obstacle is that M_f is not hyperbolic when f is reducible, and consequently convex hulls in the universal cover are not as well-behaved. Instead, we use the Bass-Serre tree T dual to the canonical reducing multicurve α for f . Suspending this canonical multicurve α in the mapping torus M_f produces the torus decomposition, and T is the tree dual to the tori. The action of $\pi_1 S$ on T thus extends to an action of $\Gamma_f = \pi_1 M_f$; see §3.1. The analogues of the hull for G and for a multicurve $u \subset \mathcal{C}(S^z)$ are then played by a G -invariant subtree $T_G \subset T$ and a K_u -invariant subtree $T_u \subset T$, respectively; see §3.3. Being purely pseudo-Anosov implies that G acts freely on T , and we show that T_G is a geometric model for G ; see Lemma 3.3. The key to proving Theorem 1.1 rests on showing that $T_G \cap T_u$ has bounded diameter, independent of $u \subset \mathcal{C}(S^z)$; see Proposition 4.1.

To understand $T_G \cap T_u$, we return to examining the convex hulls in $\mathbb{H}^2$. The splitting of $\Gamma_f \cong \pi_1 S \rtimes \langle f \rangle$ gives an action of Γ_f on $\partial\mathbb{H}^2$ by homeomorphisms, extending the isometric action of $\pi_1 S$ by deck transformations; see §3.1. This allows us to define $\mathfrak{H}_G$, which admits an isometric action by $G_0 = G \cap \pi_1 S$. Further, there is

¹It was shown in [CLM12] that admissibly embedded right-angled Artin subgroups are quite abundant in mapping class groups. See also [Run21].

a G_0 -equivariant inclusion $T_G \rightarrow \mathfrak{H}_G$, since the action on T_G is free. The quotient $p_0: \mathfrak{H}_G \rightarrow \mathfrak{H}_G/G_0 = \Sigma_0$ is an infinite type, two ended surface, and $T_G/G_0 = \sigma_0$ is a spine; see §3.4. The quotient group $G/G_0 \cong \mathbb{Z}$ admits a cocompact, non-isometric action on Σ_0 . While the action of G/G_0 on Σ_0 is not isometric, the induced action on the spine σ_0 is; see §3.5.

A critical technical step in our proof is the construction of a compact subsurface $\Sigma_1 \subset \Sigma_0$ so that for any simplex $u \subset \mathcal{C}(S^z)$, there is an element $g \in G$ for which $p_0(T_G \cap T_{g(u)}) \subset \sigma_0$ is contained in the subsurface Σ_1 ; see Lemma 4.5. We say a simplex $u \subset \mathcal{C}(S^z)$ is *deep* if $p_0(T_G \cap T_u) \subset \Sigma_1$, and by the previous sentence, it suffices to bound the diameter of $T_G \cap T_u$ for deep simplices. The construction of Σ_1 is outlined in the subsection below, but we note that $\pi_1 \Sigma_1 = G_1 < G_0$ is a finitely generated, purely pseudo-Anosov subgroup.

For a deep simplex $u \subset \mathcal{C}(S^z)$, the intersection $\mathfrak{H}_G \cap \mathfrak{H}_u$ is contained in a uniformly bounded neighborhood of $\mathfrak{H}_{G_1} \cap \mathfrak{H}_u$. Since $G_1 < \pi_1 S$ is finitely generated and purely pseudo-Anosov, this has uniformly bounded diameter from [KLS09], as described above. The vertices of the subtrees T_G and T_u are precisely those which are dual to regions of $\mathbb{H}^2 \setminus p^{-1}(\alpha)$ that $\mathfrak{H}_G$ and $\mathfrak{H}_u$ respectively intersect. If we also knew that the vertices of $T_G \cap T_u$ were dual to regions intersected by $\mathfrak{H}_G \cap \mathfrak{H}_u$, then for deep simplices, the bound on the diameter of $\mathfrak{H}_G \cap \mathfrak{H}_u$ would imply one for $T_G \cap T_u$, and we would be done. However, it is possible for both $\mathfrak{H}_G$ and $\mathfrak{H}_u$ to intersect the region dual to a vertex $t \in T$, while $\mathfrak{H}_G \cap \mathfrak{H}_u$ is disjoint from it. Our proof thus splits into two parts: bounding the diameter of a single subtree spanned by “hull type” vertices of $T_G \cap T_u$ that *do* come from $\mathfrak{H}_G \cap \mathfrak{H}_u$, and the complementary subtrees of “non-hull type” vertices that do not; see §4.2.2. The former are handled as just explained; see Lemma 4.7. For the latter, we proceed as follows.

For every path $\ell \subset T_G \cap T_u$ containing only “non-hull type” vertices, we produce geodesics $\delta_G \subseteq \partial \mathfrak{H}_G$ and $\delta_u \subseteq \partial \mathfrak{H}_u$ that “run parallel” through the regions of $\mathbb{H}^2 \setminus p^{-1}(\alpha)$ corresponding the vertices of ℓ . Hence we call the non-hull type vertices, “parallel type” vertices. If ℓ is long enough, then we show that $p_0(\delta_G) \subset \Sigma_1$ must project to a closed boundary geodesic of Σ_1 . Since this geodesic represents an element of G_1 , it is pseudo-Anosov, and so further projects to a filling geodesic in S by a result of Kra [Kra81]; see Theorem 2.2. On the other hand, δ_u projects to a simple closed geodesic in S (isotopic to a component of $\Phi(u)$; see §2.2). But if ℓ is too long, then the simple closed geodesic image of δ_u runs parallel to the filling geodesic image of δ_G for a long time, which is a contradiction. This proves a uniform bound on the diameter of ℓ ; see Lemma 4.12. Combining the hull-type and parallel-type subtree bounds for deep simplices, proves a uniform bound on $\text{diam}(T_G \cap T_u)$, for every simplex $u \subset \mathcal{C}(S^z)$, as required.

1.3. Construction of Σ_1 . We now outline the construction of the compact subsurface $\Sigma_1 \subset \Sigma_0$ with the property that for every simplex $u \subset \mathcal{C}(S^z)$, the image $p_0(T_G \cap T_u)$ can be translated into Σ_1 by an element of G . Since the spine $\sigma_0 = p_0(T_G)$ of Σ_0 has an isometric action of $\mathbb{Z} \cong G/G_0$, it suffices to prove a uniform bound on $p_0(T_G \cap T_u)$ in σ_0 : we then simply take a sufficiently large neighborhood σ_1 of a fundamental domain for the action of G/G_0 on the spine $\sigma_0 \subset \Sigma_0$, and take Σ_1 to be a thickening of σ_1 in Σ_0 .

To bound the diameter of $p_0(T_G \cap T_u)$ in σ_0 , we utilize Masur and Minsky’s subsurface projections [MM00]. For simplicity we describe the idea in the case where our reducible surface homeomorphism $f: S \rightarrow S$ is a Dehn twist about a single curve α . First, we let $A \rightarrow S$ be the annular cover whose core curve is α , and for every simplex $u \subset \mathcal{C}(S^z)$, let $\pi(v)$ be the subsurface projection of $v = \Phi(u) \subset \mathcal{C}(S)$ to the *arc graph* of A ; see §5.2. Next, for every edge $e \subset T_G$, there is a dual geodesic $\tilde{\alpha}_e \subset p^{-1}(\alpha)$, and we can identify the annulus A with the quotient $A = \mathbb{H}^2 / \text{Stab}_{\pi_1 S}(\tilde{\alpha}_e)$. There are two boundary components of $\partial \mathfrak{H}_G$ that non-trivially intersect $\tilde{\alpha}_e$, and we let Δ_e denote their image in A , viewed as a subset of the arc graph of A ; see §5.1. These sets Δ_e decorate the edges e of T_G and are G -equivariant, with $\Delta_{g(e)} = f^{\phi(g)}(\Delta_e)$ where $f^{\phi(g)}$ is the image of $g \in G$ under the homomorphism $G \rightarrow \langle f \rangle$; see Lemma 5.2.

The key idea is now the following: any edge e for which the intersection number of an arc in Δ_e with one from $\pi(v)$ is sufficiently large is a “dead end”, beyond which the hull type subtree and any parallel type subtree cannot extend; see Lemma 5.6 and Lemma 5.9. Since σ_0 has finite valence, distant vertices in σ_0 should basically differ by a large power of the generator of G/G_0 , which is isotopic to a large power of the Dehn twist f . Thus for a geodesic ℓ in $T_G \cap T_u$ with $p_0(\ell)$ sufficiently long, there must be two edges e, e' of ℓ and an element of $g \in G$ so that $e' = g(e)$ and $|\phi(g)|$ is large. However, if Δ_e and $\pi(v)$ have small intersection number (because e is not a “dead end”), then $\Delta_{e'} = \Delta_{g(e)} = f^{\phi(g)}(\Delta_e)$ and $\pi(v)$ will have large intersection number (depending on $|\phi(g)|$). Thus we get a bound on how large $|\phi(g)|$ can be, and hence a bound on how long $p_0(\ell)$ can be.

Our proof in the general case of arbitrary reducible f follows the same basic idea using the subsurface projection to the complementary components of $S \setminus \alpha$ to give a decoration on the vertices of T_G in addition to decorations of edges coming from the annular covers. The argument is complicated by the fact that f may act trivially on some subsurfaces and some annuli; see §5 for details.

Acknowledgments. The first author thanks Spencer Dowdall, Autumn Kent, and Saul Schleimer for their collaboration on [KLS09] and [DKL14], which served as motivation for the current work. Both authors thank Jason Behrstock and Dan Margalit for comments on an earlier draft of the paper. The first author was supported by NSF grant DMS-2106419 and the second by NSF grant DMS-2103191.

2. PRELIMINARIES

Throughout, S will denote a connected orientable finite type surface with negative Euler characteristic. We will equip this surface with a complete hyperbolic metric of finite area that identifies $\mathbb{H}^2$ with the universal cover $p: \mathbb{H}^2 \rightarrow S$. Given a point $z \in S$, we let S^z denote the surface obtained by puncturing S at z . We also equip S^z with a complete, finite area hyperbolic metric. The *curve complex* of S (or S^z) is the flag simplicial complex $\mathcal{C}(S)$ whose vertices are isotopy classes of essential, simple closed curves on S with two isotopy classes joined by an edge if they have disjoint representatives. Each vertex of $\mathcal{C}(S)$ has a unique geodesic representative and two vertices will be joined by an edge if and only if these geodesic representative are disjoint. Hence, each simplex of $\mathcal{C}(S)$ corresponds to a multicurve on S , which has a unique geodesic representative. Whenever convenient, we will assume that a simplex/multicurve $v \subset \mathcal{C}(S)$ is represented in S as a geodesic multicurve.

Given a surface with boundary Y , we define the *arc and curve complex* to be the flag simplicial complex $\mathcal{AC}(Y)$ whose vertices are isotopy classes of both essential, simple closed curves and isotopy classes of essential arcs meeting the boundary of Y precisely in their endpoints.² As with the curve complex, two vertices of $\mathcal{AC}(Y)$ are joined by an edge if there are disjoint representatives for the isotopy classes.

When S is a once-punctured torus or four-punctured sphere, one usually makes a different definition for $\mathcal{C}(S)$, but we do not do that here. In particular, these curve complexes are discrete, countable sets. On the other hand, if Y is a torus with one boundary component or a sphere with at least one boundary component and the sum of the number of boundary components and punctures equal to 4, then we do take the usual modified definition for $\mathcal{AC}(Y)$ in which vertices are joined by an edge if they intersect once or twice for these two types of surfaces, respectively. The reason is that for $\mathcal{C}(S)$, we need Theorem 2.3 below to hold, while for $\mathcal{AC}(Y)$, we will use coarse geometric properties in §5.

If $A \rightarrow S$ is an annular cover, let $\bar{A}$ denote the compact annulus obtained from A by adding its ideal boundary from the hyperbolic metric on S . This compactification, $\bar{A}$, of A is independent of the choice of metric. The *arc complex* $\mathcal{A}(A)$ is the flag simplicial complex whose vertices are isotopy classes of essential arcs on $\bar{A}$, where unlike other surfaces with boundary, isotopies of $\bar{A}$ are required to be the identity on $\partial\bar{A}$. Edges of $\mathcal{A}(A)$ correspond to pairs of isotopy classes with representatives having disjoint interiors. The annuli of primary interest come from curves $w \in \mathcal{C}(S)$. More precisely, every such curve w determines a conjugacy class of cyclic subgroups of $\pi_1 S$ and hence an annular covering (unique up to isomorphism) $A = A_w \rightarrow S$ for which w lifts to the core curve.

2.1. Mapping class groups and Birman exact sequence. We recall that the mapping class group of S is the group of orientation preserving homeomorphisms (or diffeomorphisms) of S , modulo the normal subgroup of those homeomorphisms that are isotopic to the identity,

$$\text{Mod}(S) = \text{Homeo}^+(S)/\text{Homeo}_0(S).$$

Every element of $\text{Mod}(S)$ is thus the isotopy class of a homeomorphism.

Recall that we have fixed a basepoint $z \in S$, and $S^z = S \setminus \{z\}$. We write $\Phi: S^z \rightarrow S$ for the inclusion. The puncture of S^z that accumulates on z via Φ is called the z -puncture and we often refer to Φ as the map that “fills the z -puncture back in”.

We are interested in the finite index subgroup $\text{Mod}(S^z, z) < \text{Mod}(S^z)$ consisting of isotopy classes of homeomorphisms that fix the z -puncture. Any homeomorphism $\varphi: S^z \rightarrow S^z$ defining an element of $\text{Mod}(S^z, z)$

²One often allows properly embedded arcs with ends in cusps of Y , if any, but we will not need such arcs in our work, so omit them in our definition.

uniquely determines a homeomorphism $\varphi': S \rightarrow S$ extending over the point z by sending it to itself and by the formula $\varphi' \circ \Phi = \Phi \circ \varphi$ on S^z . When the context makes the meaning clear, we usually abuse notation and use the same symbol φ to denote the mapping class in $\text{Mod}(S^z, z)$, a representative homeomorphism of S^z , as well as the unique extension to a homeomorphism of S .

The extension of a homeomorphism of (S^z, z) over the point z via the map Φ defines a surjective homomorphism $\Phi_*: \text{Mod}(S^z, z) \rightarrow \text{Mod}(S)$, and the *Birman's exact sequence* [Bir69] identifies an isomorphism of the kernel of Φ_* with $\pi_1 S$:

$$1 \longrightarrow \pi_1 S \longrightarrow \text{Mod}(S^z, z) \xrightarrow{\Phi_*} \text{Mod}(S) \longrightarrow 1.$$

It will be useful to describe explicitly the isomorphism of the kernel of Φ_* with $\pi_1 S$. If $\varphi: S^z \rightarrow S^z$ represents an element of the kernel, then the extension $\varphi: S \rightarrow S$ over the point z is isotopic to the identity, by an isotopy that does *not preserve* z . If $\varphi_t: S \rightarrow S$ is the isotopy so that $\varphi_0 = \varphi$ and $\varphi_1 = \mathbf{1}_S$, then defining $\gamma(t) = \varphi_t(z)$, we see that γ is a loop based at z . The isomorphism of the kernel with $\pi_1 S$ assigns the homotopy class of γ to $\varphi \in \text{Mod}(S^z, z)$. Alternatively, we can think of producing a homeomorphism $\varphi: S^z \rightarrow S^z$ by *pushing* z around the loop γ^{-1} by an isotopy on S ; we call this the *point push around* γ^{-1} .

Another perspective is useful in our setting. Fix a point $\tilde{z} \in p^{-1}(z)$. Any homeomorphism $\varphi: S^z \rightarrow S^z$ (representing an element of $\text{Mod}(S^z, z)$) has a unique lift $\tilde{\varphi}: \mathbb{H}^2 \rightarrow \mathbb{H}^2$ fixing $\tilde{z}$. The lift $\tilde{\varphi}$ is a quasi-isometry,³ and so has a unique extension to a homeomorphism $\partial\mathbb{H}^2 \rightarrow \partial\mathbb{H}^2$. Any other representative of the isotopy class of φ in $\text{Mod}(S^z, z)$ has the same extension, since the lift of the isotopy moves all points a bounded hyperbolic distance, thus we obtain an action of $\text{Mod}(S^z, z)$ on $\partial\mathbb{H}^2$.

Next observe that if $\varphi_0: S^z \rightarrow S^z$ represents an element in the kernel of Φ_* , and $\varphi_t: S \rightarrow S$ is the isotopy to the identity. This isotopy lifts to an isotopy $\tilde{\varphi}_t$ from the lift $\tilde{\varphi}_0$ fixing $\tilde{z}$ to a lift of the identity. The resulting lift of the identity, $\tilde{\varphi}_1$, is thus a covering transformation, namely the one associated to γ (as defined above by $\gamma(t) = \varphi_t(z)$). Thus, we have the following:

Proposition 2.1 (c.f. [LMS11, §1.2.3]). *The restriction of the action of $\text{Mod}(S^z, z)$ on $\partial\mathbb{H}^2$ to $\pi_1(S)$ agrees with the extension of the isometric covering action of $\pi_1 S$ on $\mathbb{H}^2$.*

Kra's Theorem [Kra81] describes precisely which elements of $\pi_1 S$ represent pseudo-Anosov elements of $\text{Mod}(S^z, z)$. Recall that a loop is *filling* if it cannot be homotoped disjoint from any essential simple closed curve (and is thus a property of the homotopy class).

Theorem 2.2 ([Kra81]). *An element of $\pi_1 S$ represents a pseudo-Anosov element of $\text{Mod}(S^z, z)$ if and only if it is represented by a filling loop.*

Since being pseudo-Anosov is equivalent to not having any isotopy classes of periodic simple closed curves, the point pushing description of Birman's isomorphism suggests a proof of Theorem 2.2; see [FM12, §14.1.4].

2.2. Fibers and trees. We let $\mathcal{C}^s(S^z) \subset \mathcal{C}(S^z)$ denote the subcomplex spanned by curves whose image under $\Phi: S^z \rightarrow S$ is an essential curve on S . We call the vertices of $\mathcal{C}^s(S^z)$ the *surviving curves* of S^z . Since Φ maps disjoint curves to disjoint curves, it induces a simplicial, surjective map which we also denote $\Phi: \mathcal{C}^s(S^z) \rightarrow \mathcal{C}(S)$, by an abuse of notation. Given any simplex, $v \in \mathcal{C}(S)$, we let $\Phi^{-1}(v)$ denote the preimage of the barycenter of v . The following is proved in [KLS09].

Theorem 2.3. *For any simplex $v \in \mathcal{C}(S)$, there is a $\pi_1 S$ -equivariant homeomorphism from the Bass-Serre tree T dual to v to $\Phi^{-1}(v) \subset \mathcal{C}^s(S^z)$. The image of a vertex $t \in T$ under this homeomorphism is the barycenter of a simplex $u_t \subset \mathcal{C}^s(S^z)$ for which $\Phi(u_t) = v$ and $\Phi|_{u_t}$ is injective. Moreover, $t, t' \in T$ are joined by an edge if and only if $u_t \cup u_{t'}$ spans a simplex of $\mathcal{C}^s(S^z)$.*

The proof of Theorem 2.3 involves some ideas that will be useful for us, which we briefly describe. Given a simplex $u \subset \mathcal{C}(S^z)$, we let K_u denote the stabilizer of u in $\pi_1 S < \text{Mod}(S^z, z)$ and $\mathfrak{H}_u \subset \mathbb{H}^2$ denote the convex hull of the limit set of K_u in $\partial\mathbb{H}^2$ (if it is nonempty). If $u \subset \mathcal{C}^s(S^z)$, $v = \Phi(u)$, and $\Phi|_u$ is injective, then $p: \mathbb{H}^2 \rightarrow S$ maps the interior $\mathfrak{H}_u^\circ \subset \mathfrak{H}_u$ to a component of $S \setminus v$ (where v is realized by its geodesic representative). Up to isotopy, $p(\mathfrak{H}_u^\circ)$ is the Φ -image of the component $U \subset S^z \setminus u$ containing the z -puncture.

³When S has cusps we assume any homeomorphism φ of S is an isometry in some neighborhood of the cusps; this is a convenience, however, as the extension of the lift to $\partial\mathbb{H}^2 \rightarrow \partial\mathbb{H}^2$ exists independent of this assumption.

One way to think about this fact is that point pushing around a loop preserves u precisely when the loop is disjoint from u , that is, when the loop (intersected with S^z) is contained in U . When $\Phi|_u$ is not injective, the component of $S^z \setminus u$ containing the z -puncture is a once-punctured annulus, making K_u an infinite cyclic group. In any case, the stabilizer of $\mathfrak{H}_u$ is exactly K_u ; see [KLS09, Theorem 4.1].

2.3. Convex cocompactness. Farb and Mosher originally defined convex compactness in the mapping class group using the action on Teichmüller space; see [FM02]. For our purposes, it will be most convenient to use the following formulation due to Kent–Leininger and independently Hamenstädt.

Theorem 2.4 ([KL08a, Ham07]). *A subgroup of the mapping class group is convex cocompact if and only if it is finitely generated and the orbit map to the curve complex is a quasi-isometric embedding.*

We will apply this to the case of subgroups of $\text{Mod}(S^z)$. We note that since the inclusion of a finite index subgroup into a bigger group is a quasi-isometry, convex cocompactness survives passage between finite index super- and subgroups.

3. SET UP

We now fix a homeomorphism $f: S \rightarrow S$ that defines an infinite order, reducible mapping class in $\text{Mod}(S)$. We let Γ denote the $\pi_1 S$ -extension group Γ_f that is the fundamental group of the mapping torus for f . Since f defines an infinite order mapping class, the homomorphism $\mu^z: \Gamma \rightarrow \text{Mod}(S^z, z)$ is injective and we identify Γ with its image in $\text{Mod}(S^z, z)$. Let $\alpha = \alpha_1 \cup \dots \cup \alpha_n \subset S$ be the canonical reduction system for the reducible mapping class defined by f ; see [Iva92]. Since convex cocompactness is preserved by passing to finite index super-groups and $\Gamma_{f^n} < \Gamma_f$ has finite index, we can replace f with a power when it is helpful. We do so, and thus (after an isotopy if necessary) assume that f fixes each curve α_i and each component of $S \setminus \alpha$. By possibly raising to a further higher power, we can also assume that f restricted to each component of $S \setminus \alpha$ is either the identity or pseudo-Anosov. We also assume throughout that α is realized as a geodesic multicurve in S with respect to our fixed hyperbolic metric.

A *complementary subsurface* of α is defined as the path metric completion Y of a component $Y^\circ \subset S \setminus \alpha$. Such a complementary subsurface Y is a hyperbolic surface with geodesic boundary and the inclusion $Y^\circ \rightarrow S \setminus \alpha$ extends to an immersion $Y \rightarrow S$ that is injective on the interior, and at most 2-to-1 at points of ∂Y . By an abuse of notation, we often write $Y \subset S$ or refer to the map $Y \rightarrow S$ as the inclusion.

Write $Y_1, \dots, Y_k$ to denote the complementary subsurfaces of α . Since each α_i and each Y_j° is invariant by f , we obtain “restricted” maps $f|_{Y_j}: Y_j \rightarrow Y_j$. We can re-index the complementary subsurfaces so that there is some $0 \leq m \leq k$ such that for $j \leq m$, $f|_{Y_j}$ is pseudo-Anosov on Y_j and for $j > m$, $f|_{Y_j}$ is the identity. We refer to these subsurfaces Y_j as the *pseudo-Anosov components* and the *identity components*, respectively.

Given f as above, we fix a finitely generated subgroup $G < \Gamma$ that is purely pseudo-Anosov as a subgroup of $\text{Mod}(S^z)$. If G is in the kernel of the homomorphism $\Phi_*: \Gamma \rightarrow \langle f \rangle < \text{Mod}(S)$, then G is contained in $\pi_1 S$, and hence is convex cocompact in $\text{Mod}(S^z)$ by [KLS09, Theorem 6.1]. Thus, we may assume that Φ_* sends G onto the subgroup of $\langle f \rangle$ generated by f^n for some $n > 0$. By passing to a further power if necessary, we can assume that $n = 1$; this will be convenient for notational purposes later and does not affect any other properties of f we have already assumed. If $\phi|_G$ is injective, then $G \cong \mathbb{Z}$, and the theorem also follows, so we assume $G_0 = \ker(\phi|_G) = G \cap \pi_1 S$ is nontrivial (hence infinite).

3.1. Action on $\mathbb{H}^2$ and T . We assume that in our fixed hyperbolic metric, the lengths of the components of α are short enough that any two components of $p^{-1}(\alpha)$ are distance at least 2 apart in $\mathbb{H}^2$. The action of $\pi_1 S$ on $\mathbb{H}^2$ preserves $p^{-1}(\alpha)$, and we let T denote Bass–Serre tree dual to $p^{-1}(\alpha)$. The action of $\pi_1 S$ on $\mathbb{H}^2$ and T extend to an action of Γ as we now explain.

As in §2.1, given any $\varphi \in \Gamma$ we write $\varphi: S^z \rightarrow S^z$ for a representative and its extension $\varphi: S \rightarrow S$ (after filling the z -puncture back in). Since $\varphi \in \Gamma$, the homeomorphism $\varphi: S \rightarrow S$ is isotopic (ignoring z) to f^k , for some $k \in \mathbb{Z}$. The lift $\tilde{\varphi}: \mathbb{H}^2 \rightarrow \mathbb{H}^2$ fixing $\tilde{z}$ is thus isotopic to a lift of f^k (not necessarily fixing $\tilde{z}$), and so has the same extension to $\partial\mathbb{H}^2$. Given any lift $\tilde{f}: \mathbb{H}^2 \rightarrow \mathbb{H}^2$ of f , any lift of f^k is then obtained by composing $\tilde{f}^k$ with an element of $\pi_1 S$. Conversely, any such composition is a lift of f^k . Hence, the action of Γ on $\partial\mathbb{H}^2$ factors through an isomorphism with the group $\langle \tilde{f}, \pi_1 S \rangle$ acting on $\partial\mathbb{H}^2$. This isomorphism $\Gamma \cong \langle \tilde{f}, \pi_1 S \rangle$ then defines an action on $\mathbb{H}^2$ extending the covering action of $\pi_1 S$. Alternatively, the given lift $\tilde{f}$ is equivariantly

isotopic to the lift $\tilde{\varphi}$ of some $\varphi \in \Gamma$ with $\Phi_*(\varphi) = f$. Then $\Gamma = \pi_1 S \rtimes \langle \varphi \rangle$ acts on $\mathbb{H}^2$ so that $\pi_1 S$ acts by covering transformations and φ^k acts by $\tilde{f}^k$ for all $k \in \mathbb{Z}$.

An alternative way to see the action of Γ on $\mathbb{H}^2$ is to consider the universal covering $\widetilde{M}_f$ of the mapping torus, lift the suspension flow, and consider the quotient by the flow lines. Since the universal cover $\mathbb{H}^2$ of S intersects each lifted flow line once, the flow space is identified with $\mathbb{H}^2$ and the action of Γ on $\widetilde{M}_f$ descends to an action of Γ on $\mathbb{H}^2$ which agrees with the covering action when restricted to $\pi_1 S$.

Since f preserves α , $\tilde{f}$ preserves $p^{-1}(\alpha)$. Therefore, Γ acts on the Bass–Serre tree T dual to α . Since f fixes each Y_i and each curve in α , a pair of vertices/edges of T are in the same Γ -orbit if and only if they are in the same $\pi_1 S$ -orbit. Unlike the action on $\mathbb{H}^2$, this action on T is by isometries. For each edge $e \subset T$, we write $\tilde{\alpha}_e$ to denote the component of $p^{-1}(\alpha)$ that is dual to e . We choose a Γ -equivariant map $\mathbb{H}^2 \rightarrow T$ sending $\tilde{\alpha}_e$ to the midpoint of e , and each component of $S \setminus p^{-1}(\alpha)$ to the $\frac{1}{2}$ -neighborhood of the dual vertex. There are many such choices, and we sometimes make a choice of one that is convenient for certain applications; for example, we may take such a map to be K -Lipschitz, where K depends only on the minimal distance between pairs of components of $p^{-1}(\alpha)$. We also choose a $\pi_1 S$ -equivariant map $T \rightarrow \Phi^{-1}(\alpha) \subset \mathcal{C}^s(S^z) \subset \mathcal{C}(S^z)$ as in Theorem 2.3, identifying vertices and edges of T with simplices of $\mathcal{C}(S^z)$ in $\Phi^{-1}(\alpha)$.

3.2. Subsurfaces and annuli for vertices and edges. For each vertex $t \in T$, we let $\tilde{Y}_t^\circ$ denote the component of $\mathbb{H}^2 \setminus p^{-1}(\alpha)$ dual to t , and use $\tilde{Y}_t$ for its closure. Let $K_t = \text{Stab}_{\pi_1 S}(\tilde{Y}_t)$ and define Y_t to be $\tilde{Y}_t/K_t$. We can identify each Y_t with exactly one of the complementary subsurfaces $Y_1, \dots, Y_k$ as follows: for a vertex $t \in T$, let $\Upsilon_t = \mathbb{H}^2/K_t$. The surface Y_t is then the convex core of Υ_t and there is a unique $i(t) \in \{1, \dots, k\}$ so that the covering map $\Upsilon_t \rightarrow S$ maps the interior of Y_t isometrically onto $Y_{i(t)}^\circ \in \{Y_1^\circ, \dots, Y_k^\circ\}$. If t and t' are in the same Γ -orbit, then Υ_t and $\Upsilon_{t'}$ are equivalent covers of S with different choices of base point. Hence there is an isomorphism of covering spaces $\Upsilon_t \rightarrow \Upsilon_{t'}$ that sends Y_t isometrically to $Y_{t'}$. In particular, $Y_{i(t)} = Y_{i(t')}$ and we use this to identify $Y_t = Y_{i(t)} = Y_{i(t')} = Y_{t'}$.

For each edge $e \subset T$, we let $K_e = \text{Stab}_{\pi_1 S}(\tilde{\alpha}_e)$ and define A_e to be the annulus $\mathbb{H}^2/K_e$. There exist a unique $i(e) \in \{1, \dots, n\}$ so that $p(\tilde{\alpha}_e) = \alpha_{i(e)}$. When convenient, we will also write $\alpha_e = \alpha_{i(e)}$. When e and e' are in the same Γ -orbit, A_e and $A_{e'}$ are equivalent annular covers of S with core curve α_e . Hence, we can isometrically identify all these annuli: $A_e = A_{i(e)} = A_{i(e')} = A_{e'}$.

We note that each vertex t and edge e of T is identified with simplices a_t and a_e of $\mathcal{C}^s(S^z)$, respectively, by Theorem 2.3, and $K_t = K_{a_t}$ and $K_e = K_{a_e}$ are indeed special cases of simplex stabilizers (so the notation is compatible with that in §2.2). Moreover, $\tilde{Y}_t = \mathfrak{H}_{a_t}$ and $\tilde{\alpha}_e = \mathfrak{H}_{a_e}$, as in §2.2. Using this, and the fact that the Γ -orbits and $\pi_1 S$ -orbits of vertices and edges of T are the same, it follows that the $\pi_1 S$ -equivariant map $T \rightarrow \Phi^{-1}(\alpha) \subset \mathcal{C}^s(S^z) \subset \mathcal{C}(S^z)$ is also Γ -equivariant.

3.3. Hulls and Trees. We now define invariant subtrees of the Bass–Serre tree T for the simplex stabilizers K_u as well as our purely pseudo-Anosov subgroup $G < \Gamma$. These subtrees will allow us to translate distances in $\mathcal{C}(S^z)$ to distances in G .

For each simplex $u \subset \mathcal{C}(S^z)$, the stabilizer $K_u < \pi_1 S$ acts by isometries on T . If the action of K_u does not have a global fixed point, we let T_u be the minimal invariant subtree of K_u . In this case, T_u is the union of the axes of loxodromic elements; see, e.g. [Bes, Proposition 2.9]. If K_u has a global fixed point in T , we define T_u to be the maximal fixed subtree. We can readily determine the structure of T_u by examining the component of $S^z \setminus u$ that contains the z -puncture.

Lemma 3.1. *Let $u \subset \mathcal{C}(S^z)$ be a multicurve and U be the component of $S^z \setminus u$ that contains the z -puncture.*

- (1) *The action of K_u on T has a global fixed point if and only if α can be isotoped to be disjoint from $\Phi(U)$ in S .*
- (2) *When K_u has a global fixed point, T_u is either a single vertex $t \in T$ or a single edge $e \subset T$. Moreover, T_u is an edge e if and only if U is a once-punctured annulus and each component of $\Phi(\partial U)$ is isotopic to the curve α_e of α .*
- (3) *If u contains a non-surviving curve, then T_u is a single vertex.*
- (4) *When u consist only of surviving curves and T_u is not an edge, then $t \in T_u$ if and only if $\mathfrak{H}_u \cap \tilde{Y}_t^\circ \neq \emptyset$.*

Proof. As described in §2.2, K_u is the group of all pushes along loops in U based at z (after filling the z -puncture back in). This group is naturally isomorphic to $\pi_1(\Phi(U), z) < \pi_1(S, z) = \pi_1 S$. Hence K_u contains a hyperbolic isometry of $\pi_1 S$ if and only if $\Phi(U)$ is not a once-punctured disk.

Now observe that $\Phi(U)$ is a once-punctured disk if and only if u contains a non-surviving curve. In this case, K_u is an infinite cyclic group generated by a parabolic isometry. Hence, there is an invariant horoball for K_u that is contained in $\tilde{Y}_t^\circ$ for some vertex $t \in T$. It follows that K_u fixes no geodesic in $p^{-1}(\alpha)$, but fixes $\tilde{Y}_t^\circ$. This implies $T_u = \{t\}$, which proves part (3).

We now focus on the case where $\Phi(U)$ is not a once-punctured disk, so K_u contains a hyperbolic isometry of $\pi_1 S$, or equivalently, when $\mathfrak{H}_u$ is non-empty.

The fixed points in T of the hyperbolic elements of $\pi_1 S$ are determined by their axes in $\mathbb{H}^2$ as follows. Let $g \in \pi_1 S$ be hyperbolic and let γ_g be the axis of g in $\mathbb{H}^2$. If $\gamma_g \subset \tilde{Y}_t^\circ$ for some vertex $t \in T$, then t is the unique fixed point of g . If $\gamma_g = \tilde{\alpha}_e$ for some geodesic $\tilde{\alpha}_e \subset p^{-1}(\alpha)$, then the edge $e \subset T$ is the maximal fixed subtree of g in T . Finally, if γ_g crosses a geodesic in $p^{-1}(\alpha)$, then periodicity says the set of geodesics in $p^{-1}(\alpha)$ that γ_g crosses will be the edges of a bi-infinite geodesic $\ell_g \subset T$. In this case, g acts loxodromically on T , and its axis in T is ℓ_g . Conversely, if g acts loxodromically on T with axis $\ell_g \subset T$, then γ_g crosses all the geodesics in $p^{-1}(\alpha)$ corresponding to the edges of ℓ_g .

Returning to K_u , if $\mathfrak{H}_u$ contains a bi-infinite geodesic that intersects a geodesic in $p^{-1}(\alpha)$ transversely, then there is some hyperbolic element of K_u whose $\mathbb{H}^2$ -axis crosses a geodesic in $p^{-1}(\alpha)$. Since this element will act loxodromically on T , K_u has no global fixed point, and thus T_u is the union of axes of the elements of K_u that act loxodromically on T . A vertex $t \in T$ is then on the T -axis of a loxodromic element $g \in K_u$ if and only if g is a hyperbolic element of $\pi_1 S$ whose $\mathbb{H}^2$ -axis intersects $\tilde{Y}_t$ in a bounded diameter segment. Thus, $t \in T_u$ if and only if $\mathfrak{H}_u \cap \tilde{Y}_t^\circ \neq \emptyset$, proving part (4). In this case, $p(\mathfrak{H}_u^\circ)$ is isotopic to $\Phi(U)$, so $\mathfrak{H}_u$ containing a geodesic that intersects a geodesic in $p^{-1}(\alpha)$ transversely ensures that α cannot be isotoped to be disjoint from $\Phi(U)$, proving (the contrapositive of) one of the implications in part (1).

If $\mathfrak{H}_u$ does not contain a geodesic that intersects a geodesic in $p^{-1}(\alpha)$ transversely, then either $\mathfrak{H}_u^\circ$ is contained entirely in $\tilde{Y}_t^\circ$ for some vertex $t \in T$, or $\mathfrak{H}_u$ is one of the geodesics $\tilde{\alpha}_e \subset p^{-1}(\alpha)$. If $\mathfrak{H}_u^\circ \subset \tilde{Y}_t^\circ$, then $T_u = \{t\}$. In this case, $p(\mathfrak{H}_u^\circ)$ is isotopic to $\Phi(U)$, so $\mathfrak{H}_u^\circ \subset \tilde{Y}_t^\circ$ implies $\Phi(U)$ is disjoint from α . If $\mathfrak{H}_u = \tilde{\alpha}_e$, then $K_u = \text{Stab}_{\pi_1 S}(\mathfrak{H}_u)$ is an infinite cyclic group generated by a hyperbolic isometry whose axis is the geodesic $\tilde{\alpha}_e$. Thus, $T_u = e$ and $p(\mathfrak{H}_u) = \alpha_e \subset \alpha$. It follows that $\Phi(\partial U)$ is an annulus with core curve isotopic to α_e . Thus, U is a once-punctured annulus and α can be isotoped to be disjoint from $\Phi(U)$. Combined with the case where $\Phi(U)$ is a once-punctured disk, this proves the other implication of part (1). Furthermore, when combined with the discussion above from the proof of part (3), we also deduce part (2). This completes the proof of the lemma. $\square$

These invariant subtrees have the following intersection property for nested simplices. This allows us to produce paths in T from paths in $\mathcal{C}(S^z)$.

Lemma 3.2. *Let u, w be simplices of $\mathcal{C}(S^z)$. If $u \subseteq w$, then $T_u \cap T_w \neq \emptyset$.*

Proof. Since $u \subseteq w$, we have $K_w < K_u$. If K_u has a global fixed point, then T_u is the maximal fixed subtree of K_u , and hence T_w is also the maximal fixed subtree of K_w . In this case, $T_u \subseteq T_w$. If neither has a global fixed point, then T_u and T_w are the minimal invariant subtrees of K_u and K_w , respectively, and so $T_w \subseteq T_u$. In either of these cases, $T_u \cap T_w \neq \emptyset$.

Finally, suppose T_u is a minimal invariant subtree for K_u and T_w is the maximal fixed subtree of K_w . By Lemma 3.1, either T_w is a vertex or edge, and in either case, there is an element $g \in K_w$ whose fixed point set is exactly T_w . Since T_u has no global fixed point, there is an axis $\ell \subset T_u$ for an element $h \in K_u$ acting loxodromically on T_u . If $\ell \cap T_w \neq \emptyset$, then $T_u \cap T_w \neq \emptyset$, as required. On the other hand, if $\ell \cap T_w = \emptyset$, then $\ell \cap g(\ell) = \emptyset$, and the geodesic from ℓ to $g(\ell)$ must non-trivially intersect T_w . Since this geodesic is contained in T_u , it follows that $T_u \cap T_w \neq \emptyset$. $\square$

Recall that we have fixed a finitely generated and purely pseudo-Anosov subgroup $G < \Gamma$ and have passed to an appropriate power of f so that $\Phi_*(G) = \langle f \rangle$. We have $G_0 = G \cap \pi_1 S$, and by our assumptions above, G_0 is an infinite, normal subgroup. Define $\mathfrak{H}_G$ to be the convex hull of the limit set of the action of G on $\partial \mathbb{H}^2$. Since we are assuming $G \neq G_0$, the action of G on $\mathbb{H}^2$ is not by isometries and does not necessarily

preserve $\mathfrak{H}_G$. However, since G_0 is a normal subgroup of G , the limit set of G_0 and G in $\partial\mathbb{H}^2$ are equal and G_0 does act isometrically on $\mathbb{H}^2$ preserving $\mathfrak{H}_G$.

Since G does act by isometries on the Bass–Serre tree T , we can use T to produce a geometric model for G as follows: since G is purely pseudo-Anosov and torsion free, no element of G fixes any simplex of $\mathcal{C}(S^z)$. Hence, G acts freely on T as its vertices and edges are Γ -equivariantly identified with simplices of $\mathcal{C}^s(S^z)$ in $\Phi^{-1}(\alpha)$. Thus, the minimal invariant subtree, T_G of the action of G on T is again the union of axes of loxodromic element of G . A compact fundamental domain for this action can be found by taking the minimal subtree containing a base vertex $v \in T_G$ and all the translates of v by a finite set of generators of G . Thus, the action of G on T_G gives G a graph of groups decomposition with trivial vertex and edge groups. This proves the following lemma.

Lemma 3.3. *The group G is free. Moreover, the tree T_G has uniformly finite valence and a free, cocompact G -action.*

Remark 3.4. Since G_0 is a normal, infinite subgroup of G , the tree T_G is also the minimal invariant tree of the action of G_0 on T . Hence T_G is also the union of the axes of the loxodromic elements of G_0 .

Since every element of G_0 is loxodromic on T_G , and since G and G_0 have equal limit sets in ∂T and $\partial\mathbb{H}^2$, a similar argument as Item (4) for Lemma 3.1 shows that the same conclusion holds for T_G and $\mathfrak{H}_G$.

Lemma 3.5. *A vertex $t \in T$ is a vertex of T_G if and only if $\tilde{Y}_t^\circ \cap \mathfrak{H}_G \neq \emptyset$.*

3.4. The G_0 -quotient and its spine. Since G_0 acts freely on T_G , there is a G_0 -equivariant embedding $T_G \rightarrow \mathfrak{H}_G$ sending vertices inside the component they are dual to (in a G -equivariant way) and sending edges to geodesic segments. Therefore, we get a surface with a spine

$$T_G/G_0 = \sigma_0 \subset \Sigma_0 = \mathfrak{H}_G/G_0.$$

Figure 1 gives an example of Σ_0 and its spine σ_0 .




FIGURE 1. Part of Σ_0 and its spine $\sigma_0 \subset \Sigma_0$. Each edge $\varepsilon \subset \sigma_0$ transversely intersects its dual arc a_ε .

Each edge $e \subset T_G$ intersects exactly one component $\tilde{\alpha}_e \subset p^{-1}(\alpha)$ and we define

$$\tilde{a}_e = \tilde{\alpha}_e \cap \mathfrak{H}_G.$$

We write $a_\varepsilon \subset \Sigma_0$ for the image of $\tilde{a}_e$ in Σ_0 where $e \subset T_G$ is an edge that projects to ε ; note that for any two edges $\varepsilon, \varepsilon'$ of σ_0 , $a_\varepsilon \cap \varepsilon'$ is empty if $\varepsilon \neq \varepsilon'$, while $a_\varepsilon \cap \varepsilon'$ is a single point if $\varepsilon = \varepsilon'$.

3.5. Polygons and G -quotient. For each vertex $t \in T_G$, the intersection $\tilde{Y}_t \cap \mathfrak{H}_G$ is an even-sided polygon with sides alternating between arcs contained in $p^{-1}(\alpha)$ and those in $\mathfrak{H}_G$. Indeed, the sides in $p^{-1}(\alpha)$ are precisely the arcs $\tilde{a}_e$ where e is an edge of T_G adjacent to t . We let $\tilde{Z}_t \subset \mathfrak{H}_G$ be this polygon corresponding to the vertex $t \in T_G$, and we write $\partial_\alpha \tilde{Z}_t$ to denote the union of the sides $\tilde{a}_e$ over all edges e adjacent to t ; see Figure 2.




FIGURE 2. Left: The polygon $\tilde{Z}_t \subset \mathfrak{H}_G \cap \tilde{Y}_t$ (shaded). Right: The “image” polygon Z_τ in Σ_0 .

Let $\tilde{p}_0: \mathbb{H}^2 \rightarrow \tilde{S}_0 = \mathbb{H}^2/G_0$ be the quotient by G_0 , which contains Σ_0 as its convex core (by definition), and write $p_0 = \tilde{p}_0|_{\mathfrak{H}_G}: \mathfrak{H}_G \rightarrow \Sigma_0$ for the restriction. Let $\eta: \tilde{S}_0 \rightarrow S$ be the associated covering corresponding to $G_0 < \pi_1 S$, so that $\eta \circ \tilde{p}_0 = p$.

Now $\eta^{-1}(\alpha) \cap \Sigma_0$ is a union of the geodesic arcs a_ε over all edges ε of σ_0 . The further restriction of p_0 to $\tilde{Z}_t$ is injective on $\tilde{Z}_t \setminus \partial_\alpha \tilde{Z}_t$ and maps $\partial_\alpha \tilde{Z}_t$ into $\eta^{-1}(\alpha)$. For a vertex $\tau \in \sigma_0$, write $Z_\tau = \tilde{Z}_t$ where t is a vertex of T_G with $p_0(t) = \tau$, and write $Z_\tau \rightarrow \Sigma_0$ to denote the restriction of p_0 . This map is injective, except possibly on the points of $\partial_\alpha \tilde{Z}_t$. As an abuse of notation, we write $Z_\tau \subset \Sigma_0$ (even though it is not necessarily embedded). See Figure 2.

Since G_0 is a normal subgroup of G , we have an action of $G/G_0 \cong \mathbb{Z}$ on $\tilde{S}_0$, and we observe that each element of G/G_0 acts as a lift of a power of f to the covering space $\tilde{S}_0$. The action of G/G_0 on $\tilde{S}_0$ is free because the action of G/G_0 on σ_0 is free. The action of $G/G_0 \cong \mathbb{Z}$ does not preserve Σ_0 , but we can find a homeomorphism $\mathfrak{f}: \tilde{S}_0 \rightarrow \tilde{S}_0$ so that $\mathfrak{f}(\Sigma_0) = \Sigma_0$ and $\mathfrak{f}$ is properly isotopic the generator of G/G_0 by an isotopy that preserves $\eta^{-1}(\alpha)$. If the generator sends a vertex $\tau \in \sigma_0$ to a vertex $\tau' \in \sigma_0$, then $\mathfrak{f}(Z_\tau) = Z_{\tau'}$; indeed, we use this to define the isotopy of the generator to the map $\mathfrak{f}$. By further proper isotopy preserving $\eta^{-1}(\alpha)$ and Σ_0 , we may assume $\mathfrak{f}(\sigma_0) = \sigma_0$. The action of $\langle \mathfrak{f} \rangle$ on Σ_0 is a topological covering space action with compact quotient Σ containing a spine $\sigma = \sigma_0 / \langle \mathfrak{f} \rangle$.

We note that the map $g \mapsto \mathfrak{f}^{\phi(g)}$ defines a homomorphism $G \rightarrow \langle \mathfrak{f} \rangle$ that descends to an isomorphism $G/G_0 \cong \langle \mathfrak{f} \rangle$. Since $\phi|_G$ is surjective, we also note that $\mathfrak{f}$ is isotopic to a lift of f . Moreover, the projection $T_G \rightarrow \sigma_0$ is equivariant with respect to this homomorphism.

4. REDUCTION TO A DIAMETER BOUND FOR $p_0(T_G \cap T_u)$

In this section we reduce the proof of Theorem 1.1 to proving that the diameter of $p_0(T_G \cap T_u)$ is uniformly bounded for all simplices $u \subset \mathcal{C}(S^z)$. This has two steps. First we show that Theorem 1.1 follows from a uniform bound on the diameter of $T_G \cap T_u$. Second we show that the diameter bound on $T_G \cap T_u$ follows from the a priori weaker bound on the diameter of $p_0(T_G \cap T_u)$. The proof of Theorem 1.1 will then be completed in §5 where we verify that $p_0(T_G \cap T_u)$ is uniformly bounded.

4.1. First reduction. We give the proof of Theorem 1.1 assuming the following proposition.

Proposition 4.1. *Given $G < \Gamma$ finitely generated and purely pseudo-Anosov in $\text{Mod}(S^z)$, there exists $D > 0$ so that for all $u \subset \mathcal{C}(S^z)$ we have*

$$\text{diam}(T_G \cap T_u) \leq D.$$

Proof of Theorem 1.1 assuming Proposition 4.1. Let $P: T \rightarrow T_G$ be the closest point projection. Observe that P maps any connected subset of $T \setminus T_G$ to a point. In particular, for any geodesic segment σ outside

T_G , $P(\sigma)$ is a point. Now suppose $u \subset \mathcal{C}(S^z)$ is any simplex and σ is a geodesic in T_u . Then $\sigma = \sigma_0 \sigma_1 \sigma_2$, where σ_1 is a (possibly empty) geodesic segment in $T_G \cap T_u$ and σ_0, σ_2 meet T_G in at most one point. It follows that $P(\sigma)$ is either a point, or $P(\sigma) = \sigma_1$. In either case, $\text{diam}(P(\sigma)) \leq D$ by Proposition 4.1.

Fix a vertex $t \in T_G$ and let $u \in \mathcal{C}(S^z)$ be a curve in the simplex that is the image of t in $\Phi^{-1}(\alpha) \subset \mathcal{C}(S^z)$. Consider the orbit map $G \rightarrow \mathcal{C}(S^z)$ given by $g \rightarrow g(u)$. Write d_T for the (geodesic) metric on T_G and d_C for the metric on the 1-skeleton of $\mathcal{C}(S^z)$.

Claim 4.2. $d_T(t, g(t)) \leq 2Dd_C(u, g(u)) + D$.

Assuming the claim, we complete the proof of the theorem. Fix a finite generating set for G and write d_G for the word metric. A standard application of the triangle inequality implies that the orbit map $G \rightarrow G \cdot u \subset \mathcal{C}(S^z)$ is lipschitz with respect to d_G . Next, note that the orbit map $G \rightarrow G \cdot t \subset T_G$ is a (κ, λ) -quasi-isometry, for some κ, λ , and thus by the claim

$$d_G(1, g) \leq \kappa d_T(t, g(t)) + \lambda \leq 2\kappa D d_C(u, g(u)) + \kappa D + \lambda.$$

Therefore, the orbit map $G \rightarrow G \cdot u \subset \mathcal{C}(S^z)$ is a quasi-isometric embedding, and hence G is convex cocompact by Theorem 2.4.

Proof of Claim 4.2. Let $n = d_C(u, g(u))$ and write $u = u_0, u_1, \dots, u_n = g(u)$ for the vertices of a $\mathcal{C}(S^z)$ -geodesic from u to $g(u)$. Consider the set of simplices

$$w_{2j} = \{u_j\} \text{ and } w_{2i+1} = \{u_i, u_{i+1}\},$$

for $j = 0, \dots, n$ and $i = 0, \dots, n-1$. In particular, $w_{2j}, w_{2j+2} \subset w_{2j+1}$ for all $j = 0, \dots, n-1$. By Lemma 3.2, this implies

$$T_{w_k} \cap T_{w_{k+1}} \neq \emptyset \text{ for all } k = 0, \dots, 2n-1.$$

We also observe that since u is a vertex of the simplex defined by t , we have $\{t\} = T_t \subseteq T_u$ and likewise $\{g(t)\} = T_{g(t)} \subseteq T_{g(u)}$.

Now construct a path $\gamma: [0, 2n+1] \rightarrow T$ by

- $\gamma(0) = t$,
- $\gamma(2n+1) = g(t)$,
- $\gamma([k, k+1]) \subset T_{w_k}$ is a geodesic segment.

This is possible because $T_{w_k} \cap T_{w_{k+1}} \neq \emptyset$ for all $k = 0, \dots, 2n-1$. Hence, we can define $\gamma(k+1)$ to be any point in the intersection of these subtrees, and then take $\gamma|_{[k, k+1]}$ to be a geodesic segment in T_{w_k} connecting the points $\gamma(k), \gamma(k+1) \in T_{w_k}$. At the endpoints, we note that $\gamma(0) = t \in T_u = T_{w_0}$ and $\gamma(2n+1) = g(t) \in T_{g(u)} = T_{w_{2n}}$.

Now consider the path $P \circ \gamma: [0, 2n+1] \rightarrow T_G$. As noted above, since

$$\gamma([k, k+1]) \subset T_{w_k},$$

Proposition 4.1 implies $\text{diam}(P \circ \gamma([k, k+1])) \leq D$. Now, $P \circ \gamma$ is a path in T_G between t and $g(t)$, and thus we have

$$d_T(t, g(t)) \leq \text{diam}(P \circ \gamma) \leq (2n+1)D = 2Dd_C(u, g(u)) + D,$$

which proves the claim. □

Having proved Claim 4.2, we have proved the theorem assuming Proposition 4.1. □

4.2. Second reduction. Having reduced the proof of Theorem 1.1 to Proposition 4.1, which asserts a uniform bound on the diameter of $T_G \cap T_u$, we proceed to our second reduction. The goal of this section is thus to deduce such a uniform bound from the following bound in $\sigma_0 = T_G/G_0$.

Proposition 4.3. *Given $G < \Gamma$ finitely generated and purely pseudo-Anosov in $\text{Mod}(S^z)$, there exists $D' > 0$ so that for any simplex $u \subset \mathcal{C}(S^z)$,*

$$\text{diam}(p_0(T_G \cap T_u)) \leq D',$$

where the diameter of $p_0(T_G \cap T_u)$ is computed in σ_0 .

We postpone the proof of Proposition 4.3 to §5 and focus this subsection on using Proposition 4.3 to prove Proposition 4.1. Our proof of Proposition 4.1 has two parts. First we show that it suffices to verify the proposition for multicurves u where $p_0(T_G \cap T_u)$ lands sufficiently deep in a specific subgraph of σ_0 . Then, we verify that $p_0(T_G \cap T_u)$ is uniformly bounded for these “deep” multicurves.

The following easy fact allows us to adjust simplices by elements of G .

Lemma 4.4. *For any $g \in G$ and simplex $u \subset \mathcal{C}(S^z)$ we have*

$$g(T_G \cap T_u) = T_G \cap T_{g(u)}.$$

Proof. Since $g \in G$, we have $g(T_G) = T_G$. Since $K_{g(u)} = gK_u g^{-1}$, by considering the two cases (minimal invariant subtree or maximal fixed subtree), we see that $g(T_u) = T_{g(u)}$. Therefore, we have

$$T_G \cap T_{g(u)} = g(T_G) \cap g(T_u) = g(T_G \cap T_u). \quad \square$$

Lemma 3.1(2) says that if T_u has finite diameter, then in fact the diameter is at most 1. Thus it will suffice to examine $T_G \cap T_u$ only for simplices $u \subset \mathcal{C}(S^z)$ where T_u has infinite diameter. In particular, by Lemma 3.1(3) we can assume all curves of u are surviving.

4.2.1. Reduction to deep simplices. Let $e_1, \dots, e_r$ be a set of $\langle f \rangle$ -orbit representatives of edges of σ_0 . Let $\sigma_1 \subset \sigma_0$ be a connected subgraph containing $e_1, \dots, e_r$ so that the distance in σ_0 from any edge e_i to a point outside σ_1 is at least $D' + 2$, where D' is the constant from Proposition 4.3. We further assume the following for each boundary component δ of $\partial\Sigma_0$: if δ^* is the minimal length loop in σ_0 that is freely homotopic in Σ_0 to δ , then there exists $i \in \mathbb{Z}$ so that $f^i(\delta^*) \subset \sigma_1$. This is possible since there are only finitely many $\langle f \rangle$ -orbits of boundary components of Σ_0 . Since σ_0 has no valence 1 vertices by virtue of being the quotient of axes of loxodromics, we can enlarge σ_1 to ensure it also has no valence 1 vertices and that σ_1 contains all edges with endpoints in its vertex set. We say that a simplex $u \subset \mathcal{C}^s(S^z)$ is *deep* if $N_2(p_0(T_G \cap T_u)) \subset \sigma_1$ where $N_2(\cdot)$ is the 2-neighborhood in σ_0 .

The next lemma, combined with Lemma 4.4, shows that it suffices to verify Proposition 4.1 for deep simplices.

Lemma 4.5. *For any simplex $u \subset \mathcal{C}(S^z)$, there exists $g \in G$ so that $g(u)$ is a deep simplex. That is,*

$$N_2(p_0(T_G \cap T_{g(u)})) \subset \sigma_1.$$

Proof. For any u , there is $j \in \mathbb{Z}$ and one of the chosen $\langle f \rangle$ -orbit representatives of edges e_i so that

$$e_i \subset f^j(p_0(T_G \cap T_u)).$$

Then $N_2(f^j(p_0(T_G \cap T_u))) \subset \sigma_1$ by Proposition 4.3.

Now we let $g \in G$ be any element that maps to f^j by the homomorphism $G \rightarrow \langle f \rangle$. Since $p_0|_{T_G}: T_G \rightarrow \sigma_0$ is equivariant with respect to this homomorphism, Lemma 4.4 implies

$$p_0(T_G \cap T_{g(u)}) = p_0(g(T_G \cap T_u)) = f^j(p_0(T_G \cap T_u)).$$

Combining this with the previous paragraph proves the lemma. $\square$

4.2.2. Subtree decomposition and bounding $T_G \cap T_u$ for deep simplices. We now use Proposition 4.3 to uniformly bound the diameter of $T_G \cap T_u$ when u is a deep simplex (and thus for any simplex, by Lemmas 4.4 and 4.5). We start by dividing the vertices of $T_G \cap T_u$ into two sets.

Definition 4.6. Given a simplex $u \subset \mathcal{C}^s(S^z)$, say that a vertex $t \in T_G \cap T_u$ is of *hull type* if

$$\mathfrak{H}_G \cap \mathfrak{H}_u \cap \tilde{Y}_t^\circ \neq \emptyset.$$

Any vertex that is not hull type is called *parallel type*.

The reason for the name “parallel type” comes from Lemma 4.10 below, which says parallel type vertices must arise from single components of $\partial\mathfrak{H}_G$ and $\partial\mathfrak{H}_u$ running parallel to each other.

The next lemma verifies that the set of hull type vertices span a subtree of $T_G \cap T_u$.

Lemma 4.7. *If the set of hull type vertices is nonempty, then it spans a subtree of $T_G \cap T_u$. That is, every vertex of the smallest subtree containing all the hull type vertices is of hull type.*

Proof. If $t, s \in T_G \cap T_u$ are hull type vertices, let $x, y \in \mathfrak{H}_G \cap \mathfrak{H}_u$ be points with $x \in \tilde{Y}_t^\circ$ and $y \in \tilde{Y}_s^\circ$. Then the geodesic $[x, y] \subset \mathbb{H}^2$ is contained in $\mathfrak{H}_G \cap \mathfrak{H}_u$ by convexity. Adjusting our equivariant map $\mathbb{H}^2 \rightarrow T$ if necessary (see §3.1), we may assume it sends $[x, y]$ to a geodesic from t to s in $T_G \cap T_u$. Every vertex of this geodesic is therefore of hull type. $\square$

We call the subtree of $T_G \cap T_u$ from Lemma 4.7 the *hull subtree*, and denote it $T_{u,G}^\mathfrak{H}$. Each maximal connected subgraph of the complement of $T_{u,G}^\mathfrak{H}$ is also a subtree of $T_G \cap T_u$. We call these components the *parallel subtrees* of $T_G \cap T_u$. To avoid arguing in separate cases, we allow the possibility that $T_{u,G}^\mathfrak{H}$ is empty (i.e. if there are no hull type vertices) in which case $T_G \cap T_u$ is the unique parallel subtree. If $T_G \cap T_u = T_{u,G}^\mathfrak{H}$, then we consider any parallel subtree to be empty.

Before bounding the diameter of the hull and parallel subtrees, we need some additional terminology. Let Σ_1 be the compact subsurface of Σ_0 defined by

$$\Sigma_1 = \bigcup_{\tau \in \sigma_1^{(0)}} Z_\tau.$$

We note that if τ, τ' are endpoints of an edge $\varepsilon \subset \sigma_1$, then there are corresponding arcs $\partial_\alpha Z_\tau, \partial_\alpha Z_{\tau'}$ which are identified in Σ_0 (hence in Σ_1) and which transversely intersect ε . Conversely, if $\tau, \tau' \in \sigma_1^{(0)}$ are vertices for which arcs of $\partial_\alpha Z_\tau$ and $\partial_\alpha Z_{\tau'}$ are identified in Σ_1 , then this arc is transverse to an edge $\varepsilon \subset \sigma_0$, which must be in σ_1 since its endpoints are. It follows that the inclusion $\sigma_1 \rightarrow \Sigma_1$ is a homotopy equivalence. Let $G_1 < G_0$ be the image of the fundamental group of Σ_1 in $G_0 = \pi_1 \Sigma_0$. Equivalently, $G_1 < G_0$ is the image of the fundamental group of σ_1 inside $G_0 = \pi_1 \sigma_0$.

Let $\tilde{\sigma}_1 \subset T_G$ be the component of $p_0^{-1}(\sigma_1)$ that is G_1 -invariant and define

$$\mathfrak{H}_G^1 = \bigcup_{t \in \tilde{\sigma}_1^{(0)}} \tilde{Z}_t.$$

Note that $\mathfrak{H}_G^1$ is the minimal, closed, G_1 -invariant subspace of $\mathfrak{H}_G$ that projects to Σ_1 . We also let $\mathfrak{H}_{G_1}$ be the convex hull of the limit set of G_1 .

Let R be the maximum of the diameters of the polygons Z_τ over all vertices $\tau \in \sigma_1$ and observe that

$$\mathfrak{H}_G^1 \subset N_R(\mathfrak{H}_{G_1}),$$

since σ_1 contains no valence 1 vertices. To see this, note that any closed loop in σ_1 without backtracking that visits every vertex of σ_1 , has geodesic representative γ in Σ_0 that meets Z_τ for every vertex $\tau \in \sigma_1$. Therefore, for every vertex $t \in \tilde{\sigma}_1$, there is a geodesic in the preimage of γ that is invariant by an infinite cyclic subgroup of G_1 and passes through $\tilde{Z}_t$. Since any such geodesic is contained in $\mathfrak{H}_{G_1}$, every point of $\mathfrak{H}_G^1$ is within R of a point of $\mathfrak{H}_{G_1}$.

We now explain how to bound the diameter of the hull subtree of $T_G \cap T_u$ for deep simplices.

Lemma 4.8. *There is a constant $D_\mathfrak{H} > 0$ so that for any deep simplex $u \subset \mathcal{C}^s(S^z)$, the diameter of $T_{u,G}^\mathfrak{H}$ is at most $D_\mathfrak{H}$.*

Proof. We may assume that T_u has infinite diameter, since otherwise it has diameter at most 1, according to Lemma 3.1(2), and the conclusion is trivial. In particular, $(\tilde{Z}_t \setminus \partial_\alpha \tilde{Z}_t) \cap \mathfrak{H}_u \neq \emptyset$ if and only if $t \in T_{u,G}^\mathfrak{H}$ by Lemma 3.1(4). Since $\mathfrak{H}_G = \bigcup_{t \in T_G^{(0)}} \tilde{Z}_t$ and $T_{u,G}^\mathfrak{H} \subset \tilde{\sigma}_1$ (because u is a deep simplex) we have

$$\mathfrak{H}_G \cap \mathfrak{H}_u = \mathfrak{H}_G^1 \cap \mathfrak{H}_u.$$

Choose our equivariant map $\mathbb{H}^2 \rightarrow T$ to be Lipschitz (see §3.1). This map sends $\mathfrak{H}_G \cap \mathfrak{H}_u$ to a set of Hausdorff distance at most $\frac{1}{2}$ from $T_{u,G}^\mathfrak{H}$, and thus it suffices to prove a bound on the diameter of the intersection $\mathfrak{H}_G \cap \mathfrak{H}_u$.

To prove such a bound, first observe that

$$\mathfrak{H}_G \cap \mathfrak{H}_u = \mathfrak{H}_G^1 \cap \mathfrak{H}_u \subset N_R(\mathfrak{H}_{G_1}) \cap \mathfrak{H}_u.$$

Now G_1 is finitely generated because Σ_1 is compact and G_1 is a purely pseudo-Anosov subgroup of $\pi_1 S < \text{Mod}(S^z, z)$ because G_0 is a purely pseudo-Anosov subgroup. The argument in §5 of [KLS09] shows that if $H < \pi_1 S < \text{Mod}(S^z, z)$ is finitely generated and purely pseudo-Anosov, then there is uniform bound on the

diameter of $N_R(\mathfrak{H}_H) \cap \mathfrak{H}_u$. In particular, there is a bound on the diameter of $N_R(\mathfrak{H}_{G_1}) \cap \mathfrak{H}_u$ determined only by G_1 , which thus also bounds the diameter of $\mathfrak{H}_G \cap \mathfrak{H}_u$. $\square$

Remark 4.9. The proof in §5 of [KLS09] actually proves a bound on $N_1(\mathfrak{H}_H) \cap \mathfrak{H}_u$, but the “1” was an arbitrary choice, and the same proof applies replacing 1 with any constant $R > 0$.

To bound the diameter of the parallel subtrees, we need the following result, which justifies the name of “parallel type” for the vertices that are not hull type.

Lemma 4.10. *Let $u \subset \mathcal{C}^s(S^z)$ be a multicurve such that T_u has infinite diameter, and let $t_0, \dots, t_n$ be the vertices of an edge path in $T_G \cap T_u$. Let e_i be the edge from t_{i-1} to t_i and $\tilde{\alpha}_i$ be the geodesic in $p^{-1}(\alpha)$ that is dual to the edge e_i . If each t_i is of parallel type, then there exists geodesics $\delta_G \subseteq \partial\mathfrak{H}_G$ and $\delta_u \subseteq \partial\mathfrak{H}_u$ so that*

- δ_u and δ_G intersect each $\tilde{\alpha}_i$ transversely;
- δ_u and δ_G do not intersect in $\tilde{Y}_{t_i}$ for any $i \in \{0, \dots, n\}$

Proof. Note, each t_i being of parallel type means that any geodesics satisfying the first item must automatically satisfy the second. Hence it suffices to produce the geodesics $\delta_G \subseteq \partial\mathfrak{H}_G$ and $\delta_u \subseteq \partial\mathfrak{H}_u$ that intersect each $\tilde{\alpha}_i$.

Convexity ensures that $\mathfrak{H}_G$ and $\mathfrak{H}_u$ intersect each $\tilde{\alpha}_i$ in a (possibly non-compact) non-empty, closed interval. If the vertices t_{i-1} and t_i are both of parallel type, then these intervals are disjoint, hence there must be $x_i \in \partial\mathfrak{H}_G \cap \tilde{\alpha}_i$ and $y_i \in \partial\mathfrak{H}_u \cap \tilde{\alpha}_i$ so that the open interval of $\tilde{\alpha}_i$ between x_i and y_i does not intersect either $\mathfrak{H}_G$ or $\mathfrak{H}_u$. Moreover, the x_i and y_i must be arranged so that the geodesic from x_i to x_{i+1} does not cross the geodesic from y_i to y_{i+1} . Let $\tilde{\alpha}_i^+ \in \partial\mathbb{H}^2$ be the endpoint of the subray of $\tilde{\alpha}_i$ starting at y_i and passing through x_i . Similarly, let $\tilde{\alpha}_i^- \in \partial\mathbb{H}^2$ be the endpoint of the subray of $\tilde{\alpha}_i$ starting at x_i and passing through y_i ; see Figure 3. Since the geodesic from x_i to x_{i+1} does not cross the geodesic from y_i to y_{i+1} , there are disjoint arcs $I_+, I_- \subset \partial\mathbb{H}^2$ so that $\tilde{\alpha}_1^+, \dots, \tilde{\alpha}_n^+ \subset I_+$ and $\tilde{\alpha}_1^-, \dots, \tilde{\alpha}_n^- \subset I_-$.

Let δ_i be the component of $\partial\mathfrak{H}_u$ that contains y_i for $i \in \{0, \dots, n-1\}$. If δ_i does not also include y_{i+1} , then δ_i must have an endpoint on the arc of $\partial\mathbb{H}^2$ between $\tilde{\alpha}_i^+$ and $\tilde{\alpha}_{i+1}^+$ (contained in I_+). But that would require δ_i to cross the geodesic from x_i to x_{i+1} as shown in Figure 3. Since the geodesic from x_i to x_{i+1}




FIGURE 3. Arrangement of x_i and y_i .

is contained in $\mathfrak{H}_G$ and δ_i is contained in $\mathfrak{H}_u$, this would contradict that t_{i+1} is a parallel type vertex for $i \in \{0, \dots, n-1\}$. Hence there is a single boundary component $\delta_u \subset \partial\mathfrak{H}_u$ that contains all the y_i . A completely analogous argument shows that there is a single boundary component $\delta_G \subseteq \partial\mathfrak{H}_G$ that contains all the x_i . $\square$

We also need this basic fact about quadrilaterals in the hyperbolic plane, the proof of which is left as an exercise in hyperbolic geometry.

Lemma 4.11. *For each $r \geq 0$ and $0 < \epsilon \leq 1$, there exists $C \geq 0$ so the following holds. Let $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ be the 4-sides of a convex quadrilateral in $\mathbb{H}^2$, labeled so that γ_1 is opposite γ_3 . If $d(\gamma_1, \gamma_3) \geq C$, then there exists subsegments $s_2 \subseteq \gamma_2$, $s_4 \subseteq \gamma_4$, each of length at least r , so that $s_2 \subseteq N_\epsilon(s_4)$ and $s_4 \subseteq N_\epsilon(s_2)$.*

We now bound the diameter of the parallel subtrees.

Lemma 4.12. *There is a constant $D_{\parallel} > 0$ so that for any deep simplex $u \subset \mathcal{C}^s(S^z)$, the diameter of any parallel subtree of $T_G \cap T_u$ is at most $D_{\parallel}$.*

Proof. By the Collar Lemma [Kee74], we may (and will) assume that the hyperbolic metric on S is chosen so that each component of α is short enough to ensure that the distance between two different geodesics in $p^{-1}(\alpha)$ is at least 1. Let $t_0, \dots, t_n$ be the vertices of a geodesic edge path in one of the parallel subtrees of $T_G \cap T_u$, then let $\tilde{\alpha}_i$ be the geodesic of $p^{-1}(\alpha)$ that is dual to the edge from t_{i-1} to t_i . By Lemma 4.10, there are geodesics $\delta_G \subseteq \partial \mathfrak{H}_G$ and $\delta_u \subseteq \partial \mathfrak{H}_u$ that form a convex quadrilateral with $\tilde{\alpha}_1$ and $\tilde{\alpha}_n$. We first show that if n is large enough, then p_0 maps δ_G onto a simple closed geodesic c that is contained in $\partial \Sigma_0 \cap \partial \Sigma_1$.

Since Σ_1 is the union of the polygons Z_τ for $\tau \in \sigma_1^{(0)}$, every component $c \subset \partial \Sigma_1$ is either a closed curve in $\partial \Sigma_0$ or $c \cap \partial \Sigma_0$ is a disjoint union of geodesic arcs. Since Σ_1 is compact, there exists $L > 0$ so that every component of $\partial \Sigma_0 \cap \partial \Sigma_1$ has length at most L .

Let δ'_G be the subsegment of δ_G between $\tilde{\alpha}_1$ and $\tilde{\alpha}_n$. Now δ'_G is the concatenation of arcs in $\partial \tilde{Z}_t \setminus \partial_\alpha \tilde{Z}_t$ where $t \in \{t_0, \dots, t_n\}$. Since the geodesics in $p^{-1}(\alpha)$ are at least 1 apart, if $n \geq L + 2$, then the length of δ'_G is at least $L + 1$. Moreover, $p_0(t_i) \in \sigma_1$ because each $t_i \in T_G \cap T_u$ and u is a deep simplex. This means $p_0(\delta'_G \cap \partial \tilde{Z}_{t_i})$ is contained in $\partial \Sigma_0 \cap \partial \Sigma_1$ for each $i \in \{1, \dots, n-1\}$, and thus $p_0(\delta'_G) \subset \partial \Sigma_0 \cap \partial \Sigma_1$. Since the components of $\partial \Sigma_0 \cap \partial \Sigma_1$ that are not closed curves in $\partial \Sigma_0$ are all arcs of length at most L , p_0 maps δ'_G onto a closed curve $c \subset \partial \Sigma_0$. Because c is a closed geodesic and δ'_G is a subsegment of δ_G , this means the entire geodesic δ_G must also map onto c .

Let $c_1, \dots, c_k$ be the closed curves in $\partial \Sigma_0 \cap \partial \Sigma_1$. For each c_i , there is a geodesic curve $\gamma_i \subset S$ so that the element of $G_1 = \pi_1 \Sigma_1 \leq \pi_1 S$ that corresponds to c_i is represented in $\pi_1 S < \text{Mod}(S^z; z)$ by the point push of z along γ_i^{-1} . Because $G_1 < G$ and G is purely pseudo-Anosov, Theorem 2.2 says each γ_i fills S , and hence these are not simple. Moreover, $\tilde{p}_0^{-1}(c_i) \subseteq p^{-1}(\gamma_i)$ where $\tilde{p}_0$ is the covering map $\mathbb{H}^2 \rightarrow \tilde{S}_0$.

The following claim puts a bound on how long a lift of a simple closed curve can travel close to a lift of one of the γ_i .

Claim 4.13. *There are $r \geq 0$ and $0 < \epsilon \leq 1$ independent of u so that for any $i \in \{1, \dots, k\}$ and any geodesic $\tilde{\gamma} \in p^{-1}(\gamma_i)$ the following holds. Let $\beta \subset S$ be a closed curve and $\tilde{\beta}$ be a geodesic in $p^{-1}(\beta)$. If $N_\epsilon(\tilde{\beta}) \cap \tilde{\gamma}$ contains a geodesic of length at least r , then β is not simple.*

Proof. Let $\epsilon \leq 1/16$ be small enough so that if x is a self intersection point of one of the γ_i , the 8ϵ -neighborhood of x on S is isometric to the 8ϵ -ball in $\mathbb{H}^2$. Let r_0 be the maximum of all the lengths of all the γ_i , then let $r = 3r_0 + 1$. These ϵ and r depend on the hyperbolic metric on S and the group G , but not on the multicurve u .

Let $\beta \subset S$ be a closed curve, then let $\tilde{\beta} \in p^{-1}(\beta)$ and $\tilde{\gamma} \in p^{-1}(\gamma_i)$ be as described in the statement of the claim. Fix a self-intersection point x of γ_i . Since $N_\epsilon(\tilde{\beta})$ contains a subsegment of $\tilde{\gamma}$ of length at least r and r is more than twice as long as the length of γ_i , there must exist $\tilde{y}_1, \tilde{w}_1, \tilde{y}_2, \tilde{w}_2 \in \tilde{\beta}$ so that the geodesic on S that connects $p(\tilde{y}_1)$ and $p(\tilde{w}_1)$ and the geodesic that connects $p(\tilde{y}_2)$ and $p(\tilde{w}_2)$ must cross; see Figure 4. Since these geodesics are subsegments of $\beta = p(\tilde{\beta})$, we have that β cannot be simple. $\square$

Let C be the constant from Lemma 4.11 for the r and ϵ from Claim 4.13. Suppose for the purposes of contradiction that $n \geq \max\{C + 2, L + 3\}$. Recall, δ_u and δ_G form a convex quadrilateral with $\tilde{\alpha}_1$ and $\tilde{\alpha}_n$. By choice of the hyperbolic metric on S , $\tilde{\alpha}_1$ and $\tilde{\alpha}_n$ are at least $n - 2 \geq C$ apart. Hence Lemma 4.11 says there is a subsegment of δ_u that is contained in the ϵ -neighborhood of δ_G . As shown above, $n \geq L + 2$ implies $p_0(\delta_G) = c_i$ for some $i \in \{1, \dots, k\}$. On the other hand, $p(\delta_u)$ is a simple curve because $p(\partial \mathfrak{H}_u) \subseteq \Phi(u)$ as described in §2.2. However, this contradicts Claim 4.13, so we must have $n < \max\{C + 2, L + 2\}$. Since C and L do not depend on u , setting $D_{\parallel} = \max\{C + 2, L + 3\}$ completes the proof of Lemma 4.12. $\square$

Armed with bounds on the diameter of the hull and parallel subtrees, we can now prove Proposition 4.1.

Proof of Proposition 4.1 assuming Proposition 4.3. Recall, we wish to prove a uniform bound D on the diameter of $T_G \cap T_u$ for every simplex $u \subset \mathcal{C}(S^z)$. We claim that setting $D = D_{\mathfrak{H}} + 2D_{\parallel} + 2$ suffices.

By Lemma 3.1 parts (2) and (3), we can assume $u \subset \mathcal{C}^s(S^z)$ and T_u has infinite diameter, while Lemmas 4.4 and 4.5 say it suffices to bound $\text{diam}(T_G \cap T_u)$ when u is a deep simplex. Let $t, t' \in T_G \cap T_u$ be any two vertices. The geodesic, ℓ , connecting t and t' decomposes into at most five segments, two contained in




FIGURE 4. A curve β cannot be simple if there is a lift $\tilde{\beta}$ that runs close to a lift $\tilde{\gamma}$ of a non-simple curve γ_i for a long enough time.

parallel subtrees, one in the hull subtree, and a pair of edges connecting the segments in parallel subtrees to the segment in the hull subtree. It follows from Lemmas 4.8 and 4.12 that the length of ℓ is at most D . Since t, t' were arbitrary, this completes the proof. $\square$

5. BOUNDING THE DIAMETER OF $p_0(T_u \cap T_G)$.

The goal of this section is to prove Proposition 4.3, which asserts the existence of a uniform bound D' on the diameter of $p_0(T_u \cap T_G)$ in σ_0 . As shown in the previous section, this will complete the proof of Proposition 4.1 and hence Theorem 1.1.

Recall from §3 that $\alpha = \alpha_1 \cup \dots \cup \alpha_n$ is the canonical reduction system of the pure, reducible homeomorphism f with complementary subsurfaces $Y_1, \dots, Y_k$. For each edge e of T (or σ_0), we write $A_e = A_{i(e)}$ for the annular cover of S corresponding to the component $\alpha_e = \alpha_{i(e)} \subset \alpha$; if e is an edge in T , then $A_e = \mathbb{H}^2/K_e$. If $e \subset T$ and $g \in \Gamma$, then we have a canonical identification $A_e = A_{g(e)}$. Likewise, for each vertex t of T (or σ_0), we write $Y_t = Y_{j(t)}$ for the corresponding complementary subsurface of S , given by $Y_t = \tilde{Y}_t/K_t$ for $t \in T$. If $t \in T$, $g \in \Gamma$, then $Y_t = Y_{g(t)}$. Observe that f acts on each $\mathcal{AC}(Y_j)$ by restricting $f|_{Y_j}$ and on each $\mathcal{A}(A_i)$ by lifting f to A_i .

We say that a vertex $t \in T$ is a *pseudo-Anosov vertex* (resp. *identity vertex*) of T if f acts by a pseudo-Anosov (resp. by the identity) on Y_t ; that is, if Y_t is a pseudo-Anosov (resp. identity) component of f . Recall, by [MM99] f will act by a pseudo-Anosov on Y_t if and only if f acts loxodromically on $\mathcal{AC}(Y_t)$. We say e is a *twist edge* of T if f acts loxodromically on $\mathcal{A}(A_e)$. This occurs if the complementary components $Y_j, Y_{j'}$ of α that meet the curve α_e are identity components, and hence f acts by a power of a Dehn twist in α_e , or if at least one of Y_j or $Y_{j'}$ are pseudo-Anosov components which effect (possibly fractional) non-canceling Dehn twists about the boundary component(s) corresponding to α_e . Since the assignments of Y_t and A_e are G -equivariant, the labeling of pseudo-Anosov/identity vertex and twist edge are also G -equivariant. Hence they descend to give the same labels to vertices and edges of $\sigma_0 = T_G/G_0$. The next lemma ensures that every path of length 2 in T contains either a twist edge or a pseudo-Anosov vertex.

Lemma 5.1. *For any non-twist edge, $e \subset T_G$, at least one endpoint is a pseudo-Anosov vertex.*

Proof. If neither endpoint of e is a pseudo-Anosov vertex, then f must act as a Dehn twist about α_e , since otherwise α_e would not be in the canonical reduction system for f . Thus, e is a twist edge. $\square$

5.1. Edge and vertex decorations. To each edge e and vertex t of T_G we will assign a bounded diameter subset Δ_e and Δ_t of the arc and curve graph of A_e and Y_t , respectively. We call these *decorations* of the edges and vertices.

For each edge e of T_G , there are exactly two geodesics in $\partial\mathfrak{H}_G$ that non-trivially intersect $\tilde{\alpha}_e$. Define $\Delta_e \subset \mathcal{A}(A_e)$ to be the union of the images of these two geodesics under the covering map $\mathbb{H}^2 \rightarrow A_e$. If e and e' are edges of T_G that are in the same G_0 -orbit, then $\Delta_e = \Delta_{e'}$ because G_0 preserves $\mathfrak{H}_G$ and $A_e = A_{e'}$.

For each vertex t in T_G , each geodesic arc $\tilde{\gamma}$ in $\tilde{Z}_t \subset \tilde{Y}_t$ with endpoints in $\partial_\alpha \tilde{Z}_t = \tilde{Z}_t \cap p^{-1}(\alpha)$ projects to a geodesic path γ in Y_t ; see Figure 5. For each such path γ , we consider the self-intersection number $\mathbb{I}(\gamma)$, which is the minimum number of double points of self intersection over all representatives of the homotopy class rel endpoints (which is realized by the unique geodesic representative orthogonal to the boundary). For each t , there are only finitely many homotopy classes of such arcs, $\gamma_1, \dots, \gamma_{r(t)}$, and we set

$$\Delta_t = \{\beta \in \mathcal{AC}(Y_t) \mid i(\beta, \gamma_j) \leq 2\mathbb{I}(\gamma_j) \text{ for some } j \in \{1, \dots, r(t)\}\}.$$

Note that by taking a representative of γ_j with only double points of self intersection realizing $\mathbb{I}(\gamma_j)$, we can construct an arc β_j in Y_t from surgery on these self intersection points, and then pushing off, so that $i(\beta_j, \gamma_j) \leq 2\mathbb{I}(\gamma_j)$. In particular, $\Delta_t \neq \emptyset$. Moreover, any β with $i(\beta, \gamma_j) \leq 2\mathbb{I}(\gamma_j)$ also has $i(\beta, \beta_j) \leq 2\mathbb{I}(\gamma_j)$ since β_j is constructed from arcs of γ_j . Since distance is bounded by a function of intersection number (see e.g. [MM99]), it follows that Δ_t has finite diameter in $\mathcal{AC}(Y_t)$. As with the edge decorations, if t and t' are vertices in the same G_0 -orbit, then $\Delta_t = \Delta_{t'}$.




FIGURE 5. Left: The polygon $\tilde{Z}_t \subset \mathfrak{H}_G$ (shaded) contained in $\tilde{Y}_t$ and essential geodesics segments contained in it. Right: The image of $\tilde{Z}_t$ and its arcs in $Y_t = \tilde{Y}_t/K_t$.

The next lemma describes how these decorations behave under arbitrary elements of G . Recall that $\phi: G \rightarrow \mathbb{Z}$ is the homomorphism so that $f^{\phi(g)} = \Phi_*(g)$ for any $g \in G$.

Lemma 5.2. *For any edge e or vertex t of T_G and $g \in G$, we have*

$$\Delta_{g(e)} = f^{\phi(g)}(\Delta_e) \quad \text{and} \quad \Delta_{g(t)} = f^{\phi(g)}(\Delta_t).$$

Proof. Observe that each $g \in G$ maps each geodesic of $\partial\mathfrak{H}_G$ to a bi-infinite path that is homotopic, rel the ideal endpoints, to a geodesic in $\partial\mathfrak{H}_G$ (since these are completely determined by the components of $p^{-1}(\alpha)$ that are intersected). Since g descends to the lift of $f^{\phi(g)}$ on each $A_e = A_{g(e)}$, the first equation follows.

For the second equation, let $\tilde{\gamma} \subset \tilde{Z}_t$ be any geodesic arc with endpoints in $\partial_\alpha \tilde{Z}_t$ and γ the image path in Y_t . Next, observe that g descends to the restriction of $f^{\phi(g)}$ to $Y_t = Y_{g(t)}$, and so maps γ to a path $f^{\phi(g)}(\gamma)$, which is homotopic to the image of a geodesic in $\tilde{Z}_{g(t)}$. Therefore, the restriction of $f^{\phi(g)}$ to Y_t maps the finite set of homotopy classes of paths defining Δ_t to those defining $\Delta_{g(t)}$, and hence sends Δ_t to $\Delta_{g(t)}$. $\square$

As a consequence, we have

Corollary 5.3. *There exists a constant $B_0 > 0$ so that*

$$\text{diam}(\Delta_e), \text{diam}(\Delta_t) \leq B_0$$

for all vertices t and edges e of T_G .

Proof. There are only finitely many G -orbits of edges and vertices in T_G and for any $g \in G$, $f^{\phi(g)}$ acts by simplicial automorphisms on $\mathcal{A}(A_e)$ and $\mathcal{AC}(Y_t)$ for every edge e and vertex t . By Lemma 5.2, it follows that

$$\text{diam}(\Delta_{g(e)}) = \text{diam}(f^{\phi(g)}(\Delta_e)) = \text{diam}(\Delta_e) \quad \text{and} \quad \text{diam}(\Delta_{g(t)}) = \text{diam}(f^{\phi(g)}(\Delta_t)) = \text{diam}(\Delta_t).$$

Therefore, we can take B_0 to be the maximum diameter of Δ_e and Δ_t taken over a finite set of G -orbit representatives of edges e and vertices t . $\square$

Since $\Delta_e = \Delta_{e'}$ and $\Delta_t = \Delta_{t'}$ for e, e' or t, t' in the same G_0 -orbit, these decorations on edges and vertices descend to decorations on the edges and vertices of $\sigma_0 = T_G/G_0$. We denote these by Δ_ε and Δ_τ for an edge ε or vertex τ of σ_0 . Since $\mathfrak{f}$ is isotopic to a lift of f to S_0 preserving Σ_0 and its spine σ_0 , the action of $G/G_0 \cong \langle \mathfrak{f} \rangle$ on $\sigma_0 = T_G/G_0$ and $\Sigma_0 = \mathfrak{H}_G/G_0$ also induces an action on the decorations, satisfying the analogous formula to Lemma 5.2:

$$(1) \quad \Delta_{\mathfrak{f}^n(\varepsilon)} = f^n(\Delta_\varepsilon) \quad \text{and} \quad \Delta_{\mathfrak{f}^n(\tau)} = f^n(\Delta_\tau)$$

for every edge ε and vertex τ of σ_0 and every $n \in \mathbb{Z}$.

5.2. Projections. Given a multicurve $v \subset \mathcal{C}(S)$, Masur and Minsky defined a projection of v to the arc and curve graph of subsurfaces and annular covers of S [MM00]. We will describe these projection in the special cases of A_e and Y_t .

For each vertex $t \in T_G$, the multicurve v intersects Y_t in a collection of disjoint curves and arcs, producing a (possibly empty) simplex of $\mathcal{AC}(Y_t)$. Let $\pi_t(v) \subset \mathcal{AC}(Y_t)$ be this simplex. We observe that $\pi_t(v)$ is precisely the set of essential arcs and curves that are in the image of $p^{-1}(v) \cap \tilde{Y}_t$ under the covering map $\tilde{Y}_t \rightarrow Y_t$ (compare with the definition of Δ_t). Since $Y_t = Y_{t'}$ if t and t' are in the same G -orbit, we have $\pi_t(v) = \pi_{t'}(v)$ in this case.

For an edge $e \in T_G$, we define $\pi_e(v) \subset \mathcal{A}(A_e)$ to be the set of essential arcs in the preimage of v under the covering map $A_e \rightarrow S$. As in the case of π_t , we note that $\pi_e(v)$ is precisely the essential arcs in the image of $p^{-1}(v)$ under the covering map $\mathbb{H}^2 \rightarrow A_e$ (compare with the definition of Δ_e). Since v is a collection of disjoint curves, $\pi_e(v)$ is a simplex of $\mathcal{A}(A_e)$. Recall, the core curve of A_e is (a lift of) one of the curves α_e in α . Thus, we have $\pi_e(v) \neq \emptyset$ if and only if $i(v, \alpha_e) \neq 0$. Since $A_e = A_{e'}$ when e and e' are in the same G -orbit, we have $\pi_e(v) = \pi_{e'}(v)$ for such pairs of edges.

Since A_e and Y_t are determined by the G -orbit of the edge or vertex, we can define projection for vertices and edges of σ_0 by

$$\pi_\varepsilon(v) = \pi_e(v) \text{ and } \pi_\tau(v) = \pi_t(v)$$

where $\varepsilon = p_0(e)$ and $\tau = p_0(t)$.

Given an edge e or vertex t of T_G (or σ_0), we let $d(\Delta_e, \pi_e(v))$ and $d(\Delta_t, \pi_t(v))$ denote the diameter of $\Delta_e \cup \pi_e(v)$ and $\Delta_t \cup \pi_t(v)$ in $\mathcal{A}(A_e)$ and $\mathcal{AC}(Y_t)$, respectively. Our proof of Proposition 4.3 hinges upon understanding for how many vertices/edges in a row these diameters can be large along a path in T_G .

For an edge $\varepsilon \in \sigma_0$ or vertex $\tau \in \sigma_0$ and $B > 0$, define:

$$\mathcal{E}(v, B) = \{\varepsilon \in \sigma_0 \mid \varepsilon \text{ is a twist edge of } \sigma_0, \pi_\varepsilon(v) \neq \emptyset, \text{ and } d(\Delta_\varepsilon, \pi_\varepsilon(v)) \leq B\}$$

$$\mathcal{V}(v, B) = \{\tau \in \sigma_0 \mid \tau \text{ is a pA vertex of } \sigma_0, \pi_\tau(v) \neq \emptyset, \text{ and } d(\Delta_\tau, \pi_\tau(v)) \leq B\}$$

We view these both as sets of edges and vertices, respectively, and as subgraphs of σ_0 (defined by taking the union of the corresponding set of edges/vertices). Let $V, E > 0$ be the numbers of G -orbits of vertices and edges in T_G , respectively. Equivalently, V, E are the numbers of $\langle \mathfrak{f} \rangle$ -orbits of vertices and edges in σ_0 , respectively.

Lemma 5.4. *For any $B > 0$ there exists $M > 0$ so that the following holds for each multicurve $v \subset \mathcal{C}(S)$:*

- (1) $\mathcal{E}(v, B)$ is a union of at most E sets of diameter at most M .
- (2) $\mathcal{V}(v, B)$ is a union of at most V sets of diameter at most M .

Proof. Fix $B \geq 0$ and a multicurve $v \subset \mathcal{C}(S)$. For part (1), it suffices to fix a twist edge $\varepsilon \subset \sigma_0$ and bound the diameter of the subset of $\mathcal{E}(v, B)$ consisting of edges in the set $\langle f \rangle \cdot \varepsilon$ by a constant M_ε , independent of v .

For this, suppose $\pi_\varepsilon(v) \neq \emptyset$. Since the monodromy f fixes the annulus A_ε , we have $\pi_{f^n(\varepsilon)}(v) = \pi_\varepsilon(v)$ for all $n \in \mathbb{Z}$. Moreover, Lemma 5.2 says $\Delta_{f^n(\varepsilon)} = f^n(\Delta_\varepsilon)$. Therefore

$$d(\Delta_{f^n(\varepsilon)}, \pi_{f^n(\varepsilon)}(v)) = d(f^n(\Delta_\varepsilon), \pi_\varepsilon(v))$$

for each $n \in \mathbb{Z}$.

Since f acts loxodromically on $\mathcal{A}(A_\varepsilon)$, the set of integers n for which $f^n(\Delta_\varepsilon)$ can intersect the B -neighborhood of $\pi_\varepsilon(v)$ is contained in an interval of integers $I_\varepsilon \subset \mathbb{Z}$ whose width, W_ε , depends only on B and the loxodromic constants of the action of f on $\mathcal{A}(A_\varepsilon)$ (and in particular, it is independent of v). Thus we have

$$\langle f \rangle \cdot \varepsilon \cap \mathcal{E}(v, B) \subseteq \bigcup_{n \in I_\varepsilon} f^n(\varepsilon).$$

The union on the right has diameter at most W_ε times the distance in σ_0 between ε and $f(\varepsilon)$ (or equivalently, the distance between $f^n(\varepsilon)$ and $f^{n+1}(\varepsilon)$, for any $n \in \mathbb{Z}$). This bound thus also bounds the diameter of $\langle f \rangle \cdot \varepsilon \cap \mathcal{E}(v, B)$, and taking any M which is at least the maximum such bound over all E orbit representatives of edges, implies part (1).

The proof for part (2) is nearly identical, choosing a pseudo-Anosov vertex and using Y_τ instead of A_ε and the fact that f acts loxodromically on $\mathcal{AC}(Y_\tau)$. $\square$

The following is an immediate corollary of Lemma 5.4 plus the bound on the valence of σ_0 (or more directly from the proof).

Corollary 5.5. *For each $B \geq 0$ there exists $N \geq 0$ so that for each multicurve $v \subset \mathcal{C}(S)$, we have*

$$|\mathcal{E}(v, B)| \leq N \quad \text{and} \quad |\mathcal{V}(v, B)| \leq N.$$

5.3. Parallel type subtrees proof. We now prove that the image under p_0 of any parallel subtree of $T_G \cap T_u$ is uniformly bounded. The main fact we need is that large vertex and edge projection can only occur along the leaves of the parallel subtrees.

Lemma 5.6. *There exists $B_1 \geq 0$ so that the following holds for each multicurve $u \subset \mathcal{C}^s(S^z)$.*

- (1) *Let ℓ be an edge path of length 2 in $T_G \cap T_u$, t be the middle vertex of ℓ , and $v = \Phi(u)$. If each vertex of ℓ is of parallel type, then $d(\Delta_t, \pi_t(v)) \leq B_1$.*
- (2) *Let ℓ be an edge path of length 3 in $T_G \cap T_u$, e be the middle edge of $T_G \cap T_u$, and $v = \Phi(u)$. If each vertex of e is of parallel type, then $d(\Delta_e, \pi_e(v)) \leq B_1$.*

Proof. First let ℓ be path of length 2 in $T_G \cap T_u$. Let e_1, e_2 be the edges of ℓ and t be the middle vertex. Let $\tilde{\alpha}_i$ be the geodesic of $p^{-1}(\alpha)$ that is dual to the edge e_i . If each vertex of ℓ is of parallel type, then Lemma 4.10 says there exist geodesics $\delta_G \subseteq \partial \mathfrak{H}_G$ and $\delta_u \subseteq \partial \mathfrak{H}_u$ that intersect $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$, but do not intersect in $\tilde{Y}_t$. Hence there is a straight line homotopy relative $\partial \tilde{Y}_t$ of $\delta_G \cap \tilde{Y}_t$ to $\delta_u \cap \tilde{Y}_t$. Since δ_u and δ_G intersect the same components of $\partial \tilde{Y}_t$, this straight line homotopy descends to a homotopy relative ∂Y_t of $p(\delta_G \cap \tilde{Y}_t)$ to $p(\delta_u \cap \tilde{Y}_t)$. In particular, $p(\delta_G \cap \tilde{Y}_t)$ is an arc on Y_t that is equal to $p(\delta_u \cap \tilde{Y}_t)$ as an element of $\mathcal{AC}(Y_t)$. Since $p(\delta_u \cap \tilde{Y}_t) \subseteq \pi_t(v)$ and $p(\delta_G \cap \tilde{Y}_t) \subseteq \Delta_t$, part (1) of the lemma now follows from Corollary 5.3 for any $B_1 \geq B_0 + 1$.

Now, let ℓ be a path of length 3 in $T_G \cap T_u$. Let e_1, e_2, e_3 be the edges of ℓ and $\tilde{\alpha}_i$ be the geodesic of $p^{-1}(\alpha)$ that is dual to the edge e_i . If each vertex of ℓ is of parallel type, then Lemma 4.10 says there exist geodesics $\delta_G \subseteq \partial \mathfrak{H}_G$ and $\delta_u \subseteq \partial \mathfrak{H}_u$ that intersect each of $\tilde{\alpha}_1, \tilde{\alpha}_2$, and $\tilde{\alpha}_3$, but do not intersect each other between $\tilde{\alpha}_1$ and $\tilde{\alpha}_3$. Now, $\tilde{\alpha}_1$ and $\tilde{\alpha}_3$ determine a unique convex ideal rectangle R ; see Figure 6. Let β_1 and β_2 be the other two sides of R .

We claim that for each non-identity element $g \in K_{e_2}$, we have $g(R) \cap R = \emptyset$. If this is true, then the rectangle R will embed onto the annulus $A_{e_2} = \mathbb{H}^2/K_{e_2}$. In particular, the images of δ_u and δ_G would be disjoint from the image of β_1 in A_{e_2} . Since $\delta_G \subseteq \partial \mathfrak{H}_G$ and $\delta_u \subseteq \partial \mathfrak{H}_u$, these project to arcs in $\Delta_{e_2} \subset \mathcal{A}(A_{e_2})$ and $\pi_{e_2}(v) \subset \mathcal{A}(A_{e_2})$, respectively, which are distance at most 2 apart in $\mathcal{A}(A_{e_2})$. Therefore,

$$d(\Delta_{e_2}, \pi_{e_2}(v)) \leq 2 + B_0 + 1 = B_0 + 3$$




FIGURE 6. The shaded rectangle R is the unique ideal rectangle with sides $\tilde{\alpha}_1$ and $\tilde{\alpha}_3$.

by Corollary 5.3, and thus, setting $B_1 = B_0 + 3$ proves part (2) of the lemma.

It therefore remains to prove that $g(R) \cap R = \emptyset$ for each non-trivial $g \in K_{e_2}$. First recall that $K_{e_2} = \text{Stab}_{\pi_1 S}(\tilde{\alpha}_2)$ is a cyclic group generated by a hyperbolic isometry of $\mathbb{H}^2$ with axis $\tilde{\alpha}_2$. Since $\tilde{\alpha}_1$ and $\tilde{\alpha}_3$ lie on different sides of $\tilde{\alpha}_2$, the only way for $g(R) \cap R \neq \emptyset$ is for $g(\tilde{\alpha}_i) \cap \tilde{\alpha}_i \neq \emptyset$ for either $i = 1$ or $i = 3$ (or both). However, no two geodesics in $p^{-1}(\alpha)$ intersect because α is a collection of disjoint simple closed curves, and every non-identity element of K_{e_2} takes every element of $p^{-1}(\alpha) - \{\tilde{\alpha}_2\}$ to a different element of $p^{-1}(\alpha) - \{\tilde{\alpha}_2\}$. Together these imply that $g(\tilde{\alpha}_i) \cap \tilde{\alpha}_i = \emptyset$ for $i \in \{1, 3\}$ and $g \in K_{e_2}$ not equal to the identity. Hence $g(R) \cap R = \emptyset$ for each non-trivial $g \in K_{e_2}$. This proves the claim above, and hence the lemma. $\square$

Combining Lemma 5.6 with Corollary 5.5 will produce the desired bound on the images of the parallel subtrees.

Lemma 5.7. *There exists $D_1 \geq 0$ so that for any multicurve $u \subset \mathcal{C}(S^z)$, if P is a parallel subtree of $T_G \cap T_u$, then $\text{diam}(p_0(P)) \leq D_1$.*

Proof. Let N be the constant from Corollary 5.5 for $B = B_1$. We claim that taking $D_1 = 3N + 6$ will suffice to prove the lemma.

Let $\tau = p_0(t)$ and $\tau' = p_0(t')$ be vertices of $p_0(P)$. Let ℓ be a geodesic path in P from t to t' and let γ be the path in σ_0 that is the image of ℓ under p_0 . Let $t = t_0, \dots, t_n = t'$ be the vertices of ℓ and let e_i be the edge of ℓ between t_{i-1} and t_i . Let $v = \Phi(u)$. By Lemma 5.6, we have

$$d(\Delta_{e_i}, \pi_{e_i}(v)) \leq B_1 \text{ and } d(\Delta_{t_j}, \pi_{t_j}(v)) \leq B_1$$

for each $i \in \{2, \dots, n-2\}$ and $j \in \{1, \dots, n-1\}$. Since

$$d(\Delta_{e_i}, \pi_{e_i}(v)) = d(\Delta_{p_0(e_i)}, \pi_{p_0(e_i)}(v)) \text{ and } d(\Delta_{t_j}, \pi_{t_j}(v)) = d(\Delta_{p_0(t_j)}, \pi_{p_0(t_j)}(v)),$$

Corollary 5.5 implies that the path $\gamma \subset \sigma_0$ contains at most $N + 2$ distinct twist edges and at most $N + 2$ distinct pseudo-Anosov vertices. Lemma 5.1 implies that every edge of γ is either a twist edge or has a pseudo-Anosov vertex as an endpoint. Hence, γ is contained in

$$\bigcup \{N_1(\tau) \mid \tau \text{ a pA vertex of } \gamma\} \cup \bigcup \{\varepsilon \mid \varepsilon \text{ a twist edge of } \gamma\}.$$

Since the diameter of $N_1(\tau)$ is at most 2, and there are at most $N + 2$ pseudo-Anosov vertices and twist edges, we have

$$\text{diam}(\gamma) \leq 2(N + 2) + (N + 2) = 3N + 6.$$

Since $\gamma = p_0(\ell)$ and ℓ is an arbitrary path in P , this implies $\text{diam}(p_0(P)) \leq 3N + 6 = D_1$ as desired. $\square$

5.4. Hull type subtree proof. Recall that $T_{u,G}^\natural \subset T_u \cap T_G$ is the hull subtree, as defined in §4.2.2. In this subsection we prove the following.

Lemma 5.8. *There exists $D_2 > 0$ so that if $u \subset \mathcal{C}^s(S^z)$ is a simplex, then $\text{diam}(p_0(T_{u,G}^\natural)) \leq D_2$.*

The first ingredient in the proof of this lemma is the following.

Lemma 5.9. *There exists a constant $B_2 \geq B_0$ with the following property. For any multicurve $u \subset \mathcal{C}^s(S^z)$, edge $e \in T_{u,G}^{\mathfrak{H}}$, and vertex $t \in T_{u,G}^{\mathfrak{H}}$, we have the following.*

- (1) *If $d(\Delta_e, \pi_e(\Phi(u))) > B_2$, then $T_{u,G}^{\mathfrak{H}} = e$.*
- (2) *If $d(\Delta_t, \pi_t(\Phi(u))) > B_2$, then t is a valence 1 vertex of $T_{u,G}^{\mathfrak{H}}$.*

Remark 5.10. Note that since $B_2 \geq B_0$, it follows that $d(\Delta_e, \pi_e(v)) > B_2$ implies $\pi_e(v) \neq \emptyset$ by Corollary 5.3. Similarly, if $d(\Delta_t, \pi_t(v)) > B_2$ then $\pi_t(v) \neq \emptyset$.

Proof. We start with the proof of part (2), which is more direct. For this, it suffices to take any $B_2 \geq B_0$. To see this, suppose $u \subset \mathcal{C}^s(S^z)$ is any multicurve, $v = \Phi(u)$, and $t \in T_{u,G}^{\mathfrak{H}}$ is a vertex with valence at least 2. This means that $\mathfrak{H}_u \cap \mathfrak{H}_G$ must intersect two distinct components $\tilde{\alpha}_e, \tilde{\alpha}_{e'} \subset p^{-1}(\alpha)$, where e, e' are adjacent to t . Let $\tilde{\gamma}$ be a geodesic arc from $\tilde{\alpha}_e$ to $\tilde{\alpha}_{e'}$ contained in $\mathfrak{H}_u \cap \mathfrak{H}_G$. Observe that $\tilde{\gamma}$ is therefore disjoint from $p^{-1}(v)$, and hence the image, γ , in Y_t is disjoint from the intersection of v with Y_t . But then, every arc and curve in $\pi_t(v)$ is disjoint from γ which implies $\pi_t(v) \subset \Delta_t$. This means that $d(\Delta_t, \pi_t(v)) \leq \text{diam}(\Delta_t) \leq B_0 \leq B_2$. Thus, any vertex of $T_{u,G}^{\mathfrak{H}}$ with $d(\Delta_t, \pi_t(v)) > B_2$ must have valence 1, proving (2).

We now explain part (1). Fix an edge $e \in T_G$ and let δ_1, δ_2 be the two geodesics in $\partial \mathfrak{H}_G$ meeting the geodesic $\tilde{\alpha}_e$ dual to e . The basic idea is that if $d(\Delta_e, \pi_e(v))$ is large, then there must be many geodesics of $p^{-1}(v)$ that cross both δ_1 and δ_2 and $\tilde{\alpha}_e$, so that for any u with $\Phi(u) = v$, a component of $\mathfrak{H}_u \cap \mathfrak{H}_G$ that meets $\tilde{\alpha}_e$ is trapped in a bounded region. We now proceed to the proof, and refer the reader to Figure 7 to aid in the argument.

First observe that there are at most four edges e' adjacent to e such that at least one of δ_1 or δ_2 crosses the geodesic $\tilde{\alpha}_{e'}$. Write $\tilde{\alpha}_i^j \subset p^{-1}(\alpha)$ for the geodesics corresponding to these edges so that δ_i non-trivially intersects $\tilde{\alpha}_i^j$ for $i, j = 1, 2$.

Next, choose an n -sheeted covering $p_e: \tilde{A}_e \rightarrow A_e$, for some $n > 0$, so that the projection $\mathbb{H}^2 \rightarrow \tilde{A}_e$ is injective on the region bounded by δ_1 and δ_2 (which contains $\mathfrak{H}_G$) as well as on each $\tilde{\alpha}_i^j$ for $i, j = 1, 2$. We denote the images of the δ_i and $\tilde{\alpha}_i^j$ in $\tilde{A}_e$ by the same name. We also use p_e to denote the induced map between arc graphs $p_e: \mathcal{A}(\tilde{A}_e) \rightarrow \mathcal{A}(A_e)$. The core curve of $\tilde{A}_e$ is an n -fold cover of the core curve α_e of A_e , and we denote it by α_e^n . See Figure 7. Observe that the degree n necessary to arrange that all of these things happen can be chosen to depend only on the G -orbit of e , and since there are only finitely many G -orbits of edges in T_G , we can in fact assume that n is independent of e .




FIGURE 7. A schematic of $\tilde{A}_e$.

Next, observe that for any arc $\gamma \in \mathcal{A}(\tilde{A}_e)$, we have

$$i(p_e(\gamma), p_e(\delta_i)) \leq n(i(\gamma, \delta_i) + 1)$$

for $i = 1, 2$. Since distances in arc graphs of annuli are given by intersection number plus 1, it follows that

$$d(p_e(\gamma), p_e(\delta_i)) \leq n d(\gamma, \delta_i) + 1,$$

for $i = 1, 2$. In particular, note that

$$\begin{aligned} d(\pi_e(v), \Delta_e) &= d(\pi_e(v), p_e(\delta_1) \cup p_e(\delta_2)) \\ &\leq n d(p_e^{-1}(\pi_e(v)), \delta_1 \cup \delta_2) + 1 \\ &\leq n(d(p_e^{-1}(\pi_e(v)), \delta_i) + 1) + 1, \end{aligned}$$

for each $i = 1, 2$, since $d(\delta_1, \delta_2) = 1$ in $\mathcal{A}(\tilde{A}_e)$.

We set $B_2 \geq \max\{B_0, 10n+1\}$. Suppose $u \subset \mathcal{C}^s(S^z)$ is any multicurve for which $T_{u,G}^\mathfrak{H}$ contains e , $v = \Phi(u)$, and $d(\pi_e(v), \Delta_e) > B_2$. Since the diameter of $p_e^{-1}(\pi_e(v))$ is 1, it follows that for any $\tilde{v} \in p_e^{-1}(\pi_e(v))$ and $i = 1, 2$, we have

$$d(\tilde{v}, \delta_i) \geq d(p_e^{-1}(\pi_e(v)), \delta_i) - 1.$$

Combining this with the inequalities above gives

$$d(\tilde{v}, \delta_i) \geq \frac{1}{n}(d(\pi_e(v), \Delta_e) - 1) - 2 > \frac{1}{n}(B_2 - 1) - 2 \geq \frac{1}{n}10n - 2 \geq 8.$$

Thus, any $\tilde{v} \in p_e^{-1}(\pi_e(v))$ intersects each of δ_1 and δ_2 in at least 7 points in $\tilde{A}_e$. Since $\tilde{v}$ and δ_i are geodesics in $\tilde{A}_e$, the difference in the number of intersection points on the two sides of the core geodesic α_e^n is at most 1. It follows that there are at least 3 points of intersection of $\tilde{v}$ with each of δ_1 and δ_2 on either side of α_e^n .

We now see that for any $\tilde{v} \in p_e^{-1}(\pi_e(v))$ there are arcs of intersection of $\tilde{v}$ in $\tilde{A}_e$ with the region bounded by δ_1 and δ_2 that contains $\mathfrak{H}_G$, on both sides of α_e^n . Moreover, there are such segments that meet δ_1 and δ_2 between the geodesics $\{\tilde{\alpha}_i^j\}_{i,j}$ (since once $\tilde{v}$ meets $\tilde{\alpha}_i^j$ it can intersect δ_i in at most one more point), and therefore, each segment is contained in the image of the corresponding Z_t , for t an endpoint of e .

Since $T_{u,G}^\mathfrak{H}$ contains e , the projection of $\mathfrak{H}_u \cap \mathfrak{H}_G$ to $\tilde{A}_e$ necessarily intersects α_e^n . It is therefore contained in the region between two of the segments of any $\tilde{v}$ described above. In fact, it follows that there are $\tilde{v}_1, \tilde{v}_2 \in p_e^{-1}(\pi_e(v))$ such that $\mathfrak{H}_u \cap \mathfrak{H}_G$ is contained in the region bounded by these two geodesics together with δ_1 and δ_2 ; see Figure 7. This implies that $T_{u,G}^\mathfrak{H} = e$, as required. This completes the proof of (1), and hence the lemma. $\square$

The second ingredient is the following bound on paths that are not contained in the subsets $\mathcal{E}(v, B)$ and $\mathcal{V}(v, B)$ from §5.1. Recall that E denotes the number of G/G_0 -orbits of edges in σ_0 .

Lemma 5.11. *Let B_2 be as in Lemma 5.9 and suppose $u \subset \mathcal{C}^s(S^z)$ is any multicurve, $v = \Phi(u)$, and $\gamma \subset T_{u,G}^\mathfrak{H}$ is an embedded edge path such that $p_0(\gamma) \subset \sigma_0$ is disjoint from $\mathcal{E}(v, B_2) \cup \mathcal{V}(v, B_2)$. Then the length of γ is at most $2E + 2$.*

Proof. Suppose that the length of γ is greater than $2E + 2$ and let $\gamma_0 \subset \gamma$ be the subpath obtained by deleting the first and last edge. Note this path contains no valence 1 vertices of $T_{u,G}^\mathfrak{H}$, and consequently by Lemma 5.9, every vertex t of γ_0 is either an identity (non-pseudo-Anosov) vertex or has $\pi_t(v) = \emptyset$. Since $T_{u,G}^\mathfrak{H}$ cannot be a single edge, another application of Lemma 5.9 implies that for every edge e of γ_0 , either e is a non-twist edge or $\pi_e(v) = \emptyset$. In fact, we claim that something stronger holds for edges.

Claim 5.12. *For every edge e of γ_0 , $\pi_e(v) = \emptyset$.*

Proof. If e is a twist edge, then we have already noted that $\pi_e(v) = \emptyset$, so it suffices to assume e is a non-twist edge. In this case, at least one of its endpoints, call it t , is a pseudo-Anosov vertex and hence $\pi_t(v) = \emptyset$. Since v cannot intersect the core curve $\alpha_e \subset A_e$ without intersecting Y_t , it follows that $\pi_e(v) = \emptyset$. $\square$

Since the length of γ_0 is greater than $2E$, there must be a pair of edges e_0, e_1 of γ_0 so that e_0 and e_1 —viewed as oriented edges, oriented by an orientation on γ_0 —differ by an element $g \in G$. Without loss of generality, suppose e_0 is the first of these edges encountered along γ_0 . Recall (see §3.4) that we have embedded T_G into $\mathfrak{H}_G \subset \mathbb{H}^2$, G -equivariantly on the vertices. Using this, we let γ_1 be the subsegment of γ_0 that begins with e_0 and ends with e_1 . Let $\nu_1 \subset \gamma_1$ be the subpath starting from $e_0 \cap \tilde{\alpha}_{e_0}$ and ending at $e_1 \cap \tilde{\alpha}_{e_1}$.

Claim 5.13. *The path ν_1 is homotopic, rel endpoints, in $\mathbb{H}^2$ to a path ν'_1 which is disjoint from $p^{-1}(v)$.*

Proof. Since $\pi_e(v) = \emptyset$ for every edge e of γ_0 , it follows that $p^{-1}(v)$ is disjoint from $\tilde{\alpha}_e$ for every such edge e . Let $t \in \nu_1$ be any vertex and let $e, e' \subset \gamma_1$ be the edges for which $t = e \cap e'$ with e appearing before e' . Let $\nu_t \subset \nu_1$ be the subpath from $e \cap \tilde{\alpha}_e$ to $e' \cap \tilde{\alpha}_{e'}$. Either we can homotope ν_t , rel endpoints, to a path disjoint from $p^{-1}(v)$, or else some component of $\tilde{v} \subset p^{-1}(v)$ separates $\tilde{\alpha}_e$ from $\tilde{\alpha}_{e'}$. The latter situation cannot happen, however, because then $\mathfrak{H}_u \cap \mathfrak{H}_G$ would have to lie on one side or the other of $\tilde{v}$, contradicting the fact that e and e' are in $T_{u,G}^\mathfrak{H}$. We can then combine the homotopies for each subpath ν_t associated to each vertex t of ν_1 , producing the required homotopy for ν_1 . $\square$

The rest of the proof splits into two cases.

Case 1. There is some vertex t in ν_1 for which $\pi_t(v) \neq \emptyset$.

Note that $\pi_t(v) \neq \emptyset$ implies t is an identity vertex. If $\pi_t(v)$ contains a simple closed curve, then denote it by w and note that it is a component of v . Otherwise, $\pi_t(v)$ is a collection of arcs with endpoints on the boundary of Y_t . In this case, let a be an arc of $\pi_t(v)$ and $c \subset \partial Y_t \subset \alpha$ be the component(s) of ∂Y_t containing the endpoints of a . The boundary of a small neighborhood of $a \cup c$ contains either one or two essential curves on Y_t ; see [MM00, §2]. Define w to be one of these curves.

In either case, by further homotopy if necessary, we may assume that ν_1' is disjoint from $p^{-1}(w)$. This is obvious if $w \subset v$, while for the other case, we argue as follows. Since $\pi_e(v) = \emptyset$ for every edge e of γ_0 , no component of c is in $p(\tilde{\alpha}_e)$ for any edge e of γ_0 . Hence, $p(\nu_1')$ is disjoint from both c and the arc a , ensuring ν_1' is disjoint from $p^{-1}(w)$.

Now, the element $g \in G$ maps $\tilde{\alpha}_{e_0}$ to $\tilde{\alpha}_{e_1}$. Let ν_1'' be the path obtained by concatenating ν_1' with an arc of $\tilde{\alpha}_{e_1}$ from the terminal endpoint of ν_1' to the g -image of the initial endpoint. Since $\tilde{\alpha}_{e_1}$ is disjoint from $p^{-1}(w)$, it follows that ν_1'' is disjoint from $p^{-1}(w)$. Now set

$$\tilde{v} = \bigcup_{n \in \mathbb{Z}} g^n(\nu_1''),$$

which is a bi-infinite, g -invariant path. Furthermore, since w is a simple closed curve contained in an identity complementary region, $p^{-1}(w)$ is invariant by g as well. In particular, $\tilde{v}$ is also disjoint from $p^{-1}(w)$. That is, $\tilde{v}$ is contained in a single component of $\mathbb{H}^2 \setminus p^{-1}(w)$, and therefore, g is contained in the stabilizer of this set. The closure of this component is $\mathfrak{H}_{u_0}$ for some curve u_0 with $\Phi(u_0) = w$, and therefore g fixes u_0 , contradicting the fact that G is purely pseudo-Anosov. This contradiction shows that Case 1 cannot happen.

Case 2. For every vertex t of ν_1 we have $\pi_t(v) = \emptyset$.

Under these assumptions, we note that v is disjoint from every complementary subsurface Y_j that $p(\nu_1)$ intersects. In particular, there must be some curve $\alpha_i \in \alpha$ that is disjoint from $p(\nu_1)$. We can then build ν_1'' and $\tilde{v}$ as we did above, but in this case, the bi-infinite path $\tilde{v}$ is disjoint from $p^{-1}(\alpha_i)$ instead of $p^{-1}(w)$. Since $p^{-1}(\alpha_i)$ is invariant by g , we again find that $\tilde{v}$ is contained in a set $\mathfrak{H}_{u_0}$ where u_0 is a curve with $\Phi(u_0) = \alpha_i$. As before, this implies g fixes u_0 , which is another contradiction. Therefore, Case 2 cannot happen either. Since these two cases account for all possibilities, we see that the assumption that γ had length greater than $2E + 2$ was impossible. $\square$

The lemma above uniformly bounds the length of any subsegment of $p_0(T_{u,G}^\mathfrak{H})$ that is outside of the set $\mathcal{E}(v, B_2) \cup \mathcal{V}(v, B_2)$. Combining this with the fact that $\mathcal{E}(v, B_2)$ and $\mathcal{V}(v, B_2)$ are finite collections of uniformly bounded diameter sets (Lemma 5.4), we can produce a uniform bound of $\text{diam}(p_0(T_{u,G}^\mathfrak{H}))$.

Proof of Lemma 5.8. Recall that E and V respectively denote the number of G/G_0 -orbits of edges and vertices in σ_0 . Let $M > 0$ be the constant from Lemma 5.4 for $B = B_2$ and set $D_2 = (E + V)(2M + 2E + 3)$.

Recall that $\mathcal{E}(v, B_2) \cup \mathcal{V}(v, B_2) \subset \sigma_0$ is a union of at most $E + V$ sets of diameter at most M by Lemma 5.4. If $T_{u,G}^\mathfrak{H}$ does not intersect any of these sets, then Lemma 5.11 says $\text{diam}(p_0(T_{u,G}^\mathfrak{H})) \leq 2E + 2 \leq D_2$. Otherwise, let $L \leq E + V$ be the number of these sets that non-trivially intersect $p_0(T_{u,G}^\mathfrak{H})$, and let $X_1, \dots, X_L$ be these sets (whose diameters are thus at most M). According to Lemma 5.11, the maximal length of an edge path in $p_0(T_{u,G}^\mathfrak{H})$ outside the union of these sets is at most $2E + 2$. Therefore the maximum distance from any point of $p_0(T_{u,G}^\mathfrak{H})$ to $X_1 \cup \dots \cup X_L$ is at most $2E + 3$. The collection $\{N_{M+2E+3}(X_j)\}_{j=1}^L$ is then a connected

cover of $p_0(T_{u,G}^5)$, where each set has diameter at most $2M + 2E + 3$. Thus, the diameter of $p_0(T_{u,G}^5)$ is at most $L(2M + 2E + 3) \leq (E + V)(2M + 2E + 3) = D_2$, as required. $\square$

5.5. Combining bounds. The proof of Proposition 4.3 is now straightforward.

Proposition 4.3. *Given $G < \Gamma$ finitely generated and purely pseudo-Anosov in $\text{Mod}(S^z)$, there exists $D' > 0$ so that for any simplex $u \in \mathcal{C}(S^z)$,*

$$\text{diam}(p_0(T_G \cap T_u)) \leq D',$$

where the diameter of $p_0(T_G \cap T_u)$ is computed in σ_0 .

Proof. Observe that $T_u \cap T_G$ is a union of $T_{u,G}^5$ and some set of parallel type subtrees, each of which is connected by an edge to $T_{u,G}^5$. Therefore, by Lemmas 5.7 and 5.8, the diameter of $p_0(T_u \cap T_G)$ is at most $2D_1 + D_2 + 2$. $\square$

REFERENCES

- [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.
- [Bes] Mladen Bestvina. Real trees in topology, geometry, and group theory, Preprint. [arXiv:math/9712210](#).
- [Bir69] Joan S. Birman. Mapping class groups and their relationship to braid groups. *Comm. Pure Appl. Math.*, 22:213–238, 1969.
- [CLM12] Matt T. Clay, Christopher J. Leininger, and Johanna Mangahas. The geometry of right-angled Artin subgroups of mapping class groups. *Groups Geom. Dyn.*, 6(2):249–278, 2012.
- [DKL14] Spencer Dowdall, Autumn E. Kent, and Christopher J. Leininger. Pseudo-Anosov subgroups of fibered 3-manifold groups. *Groups Geom. Dyn.*, 8(4):1247–1282, 2014.
- [DT15] Matthew Gentry Durham and Samuel J. Taylor. Convex cocompactness and stability in mapping class groups. *Algebr. Geom. Topol.*, 15(5):2839–2859, 2015.
- [FLP91] A. Fathi, F. Laudenbach, and V. Poénaru. *Travaux de Thurston sur les surfaces*. Société Mathématique de France, Paris, 1991. Séminaire Orsay, Reprint of *Travaux de Thurston sur les surfaces*, Soc. Math. France, Paris, 1979 Astérisque No. 66-67 (1991).
- [FM02] B. Farb and L. Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152 (electronic), 2002.
- [FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.
- [Ham] Ursula Hamenstädt. Word hyperbolic extensions of surface groups. Preprint, [arXiv:math.GT/0505244](#).
- [Ham07] Ursula Hamenstädt. Geometry of the complex of curves and of Teichmüller space. In *Handbook of Teichmüller theory. Vol. I*, volume 11 of *IRMA Lect. Math. Theor. Phys.*, pages 447–467. Eur. Math. Soc., Zürich, 2007.
- [Iva92] N. V. Ivanov. *Subgroups of Teichmüller modular groups*, volume 115 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1992. Translated from the Russian by E. J. F. Primrose and revised by the author.
- [Kee74] Linda Keen. Collars on Riemann surfaces. In *Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973)*, pages 263–268. Ann. of Math. Studies, No. 79. Princeton Univ. Press, Princeton, N.J., 1974.
- [KL07] Autumn E. Kent and Christopher J. Leininger. Subgroups of mapping class groups from the geometrical viewpoint. In *In the tradition of Ahlfors-Bers. IV*, volume 432 of *Contemp. Math.*, pages 119–141. Amer. Math. Soc., Providence, RI, 2007.
- [KL08a] Autumn E. Kent and Christopher J. Leininger. Shadows of mapping class groups: capturing convex cocompactness. *Geom. Funct. Anal.*, 18(4):1270–1325, 2008.
- [KL08b] Autumn E. Kent and Christopher J. Leininger. Uniform convergence in the mapping class group. *Ergodic Theory Dynam. Systems*, 28(4):1177–1195, 2008.
- [KLS09] Autumn E. Kent, Christopher J. Leininger, and Saul Schleimer. Trees and mapping class groups. *J. Reine Angew. Math.*, 637:1–21, 2009.
- [KMT17] Thomas Koberda, Johanna Mangahas, and Samuel J. Taylor. The geometry of purely loxodromic subgroups of right-angled Artin groups. *Trans. Amer. Math. Soc.*, 369(11):8179–8208, 2017.
- [Kra81] I. Kra. On the Nielsen-Thurston-Bers type of some self-maps of Riemann surfaces. *Acta Math.*, 146(3-4):231–270, 1981.
- [LMS11] Christopher J. Leininger, Mahan Mj, and Saul Schleimer. The universal Cannon-Thurston map and the boundary of the curve complex. *Comment. Math. Helv.*, 86(4):769–816, 2011.
- [MM99] H. A. Masur and Y. N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999.

- [MM00] H. A. Masur and Y. N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.
- [MS12] Mahan Mj and Pranab Sardar. A combination theorem for metric bundles. *Geom. Funct. Anal.*, 22(6):1636–1707, 2012.
- [Run21] Ian Runnells. Effective generation of right-angled Artin groups in mapping class groups. *Geom. Dedicata*, 214:277–294, 2021.
- [Tsh21] Bena Tshishiku. Convex cocompact subgroups the goeritz group. Preprint, [arXiv:2109.06981](https://arxiv.org/abs/2109.06981), 2021.

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

DEPARTMENT OF MATHEMATICS, RICE UNIVERSITY, HOUSTON, TX
Email address: jr92@rice.edu