COMPOSITIO MATHEMATICA

Right-angled Artin groups as normal subgroups of mapping class groups

Matt Clay, Johanna Mangahas and Dan Margalit

Compositio Math. **157** (2021), 1807–1852.

doi:10.1112/S0010437X21007417

Compositio Math. 157 (2021) 1807–1852
doi:10.1112/S0010437X21007417

Right-angled Artin groups as normal subgroups of mapping class groups

Matt Clay, Johanna Mangahas and Dan Margalit

Abstract

We construct the first examples of normal subgroups of mapping class groups that are isomorphic to non-free right-angled Artin groups. Our construction also gives normal, non-free right-angled Artin subgroups of other groups, such as braid groups and pure braid groups, as well as many subgroups of the mapping class group, such as the Torelli subgroup. Our work recovers and generalizes the seminal result of Dahmani–Guirardel–Osin, which gives free, purely pseudo-Anosov normal subgroups of mapping class groups. We give two applications of our methods: (1) we produce an explicit proper normal subgroup of the mapping class group that is not contained in any level m congruence subgroup and (2) we produce an explicit example of a pseudo-Anosov mapping class with the property that all of its even powers have free normal closure and its odd powers normally generate the entire mapping class group. The technical theorem at the heart of our work is a new version of the windmill apparatus of Dahmani–Guirardel–Osin, which is tailored to the setting of group actions on the projection complexes of Bestvina–Bromberg–Fujiwara.

1. Introduction

This paper is an investigation into the structure of normal subgroups of the mapping class group. Although there is no hope of a complete classification of such subgroups into isomorphism types, we may hope for broader descriptions of the various possible behaviors. One of the main goals of this paper is to give new examples of right-angled Artin groups (RAAGs) that embed as normal subgroups of mapping class groups.

1.1 Overview

We begin with an overview of the content, and context, of this paper. The first examples of normal, right-angled Artin subgroups of mapping class groups of arbitrary surfaces were given in the celebrated work of Dahmani–Guirardel–Osin [DGO17]. They proved that if f is pseudo-Anosov, then the normal closure of some high power of f is a free group of infinite rank.

The Dahmani–Guirardel–Osin result cannot be generalized to arbitrary mapping classes of infinite order. Indeed, if f is a mapping class with sufficiently small support (say, a power of a

Received 2 March 2020, accepted in final form 12 February 2021, published online 27 July 2021.
2020 Mathematics Subject Classification 57K20 (primary), 57M07, 20F65 (secondary).
Keywords: mapping class group, right-angled Artin group.

Dehn twist about a non-separating curve), then f has a conjugate that commutes with it. Thus, the normal closure of f is not free.

In the absence of freeness, we may hope that the normal closure of a power of f is isomorphic to a RAAG, that is, a group where all of the defining relations are commutations among generators. However, Brendle and the third author [BM19] showed that if the support of a mapping class is sufficiently small (in a precise sense that they define) then its normal closure is not isomorphic, or even abstractly commensurable, to a RAAG; see also [CLM14].

In summary: if the support of a mapping class is the entire surface (the pseudo-Anosov case), then a large power has normal closure a free group, and if the support is sufficiently small, then the normal closure cannot be any RAAG. The main result of this paper, Theorem 1.1 below, can be summarized as follows:

If the support of a mapping class is sufficiently large, then it has a large power whose normal closure is a right-angled Artin group.

As we show, Theorem 1.1 applies not only to single elements but also to finite collections. Theorem 1.1 further applies to normal closures in arbitrary subgroups of the mapping class group as well. Thus, for example, we may apply Theorem 1.1 in order to construct new normal, right-angled Artin subgroups of the pure braid group and the Torelli group.

Theorem 1.1 gives the precise isomorphism types of the RAAGs that arise from our construction. Specifically, each is a free product of groups of the following form:

$$F_\infty, \quad \mathop{\ast}_{\infty}(F_\infty \times F_\infty), \quad \mathop{\ast}_{\infty}(F_\infty \times \mathbb{Z}), \quad \text{and} \quad \mathop{\ast}_{\infty}(F_\infty \times F_\infty \times \mathbb{Z}).$$

The 14 non-free RAAGs arising as a free product in this way are the first known examples of non-free, normal, right-angled Artin subgroups of the mapping class group. We show in § 2 that all 15 RAAGs described previously (including the free RAAG) also appear as normal subgroups of the Torelli group and the pure braid group.

As applications of Theorem 1.1 we exhibit two new phenomena.

(i) There is a normal subgroup of the mapping class group that is not contained in any level m congruence subgroup.
(ii) There is a pseudo-Anosov mapping class with the property that all of its odd powers have normal closure equal to the mapping class group and all of its non-zero even powers have normal closure a free group of infinite rank.

See Theorems 1.2 and 1.5 for the precise statements. The former answers a question raised in earlier work by Lanier and the third author.

In order to prove Theorem 1.1 we appeal to, and develop, the theory of projection complexes, introduced by Bestvina–Bromberg–Fujiwara [BBF15]. We give a general result (Theorem 1.6), which says that if a group has a 'spinning' action on a projection complex, then the group is isomorphic to a free product of certain vertex stabilizers. The projection complexes that we consider in the proof of Theorem 1.1 are novel in that the vertices correspond to disconnected subsurfaces of the given surface. Theorem 1.6 applies much more generally, though, and has applications, for example, to the theory of $\mathrm{Out}(F_n)$ (see § 1.7).

Outline of the introduction. The rest of the introduction is structured as follows. We begin by giving the statement of our main result in § 1.2. We give some first consequences in § 1.3, where

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

we explicitly construct right-angled Artin subgroups of the mapping class group by directly applying Theorem 1.1. In § 1.4 we discuss Theorem 1.2, our construction of normal subgroups of the mapping class group that are not contained in congruence subgroups. In § 1.5 we give a complete picture of which RAAGs arise from our construction (the 15 examples discussed previously). Then in § 1.6 we discuss some of the finer details about the statement of our main theorem, addressing the question: given a pseudo-Anosov mapping class, precisely which powers do, and do not, have free normal closure? Here we state Theorem 1.5. In § 1.7 we discuss our technical results about group actions on projection complexes. Finally, in § 1.8 we give an outline of the remainder of the paper.

1.2 Statement of the main result

The statement of our main result, Theorem 1.1 below, requires a number of definitions and some notation.

Mapping class groups and pseudo-Anosov mapping classes. Let S_g denote the closed, connected, orientable surface of genus g, and let $S_{g,p}$ denote the surface obtained from S_g by deleting p points (so $S_{g,0} = S_g$). Finally, let $S^b_{g,p}$ be the surface obtained from $S_{g,p}$ by deleting the interiors of b disjoint disks (so $S^0_{g,p} = S_{g,p}$). The mapping class group $\mathrm{Mod}(S^b_{g,p})$ is defined as the group of connected components of $\mathrm{Homeo}^+(S^b_{g,p}, \partial S^b_{g,p})$, the group of orientation-preserving homeomorphisms of $S^b_{g,p}$ that fix the boundary pointwise. In what follows, fix $S = S^b_{g,p}$.

The Nielsen–Thurston classification theorem states that each element of $\mathrm{Mod}(S)$ is either periodic, reducible, or pseudo-Anosov; see, e.g., [FM12, Chapter 13]. The group $\mathrm{Mod}(S)$ acts on the space of projective measured foliations $\mathrm{PMF}(S)$ and an element is pseudo-Anosov if and only if the cyclic subgroup it generates acts with source–sink dynamics; in this case the source is denoted F_- and the sink F_+.

Partial pseudo-Anosov mapping classes. By a subsurface of S we will mean a closed submanifold X with the property that each component of ∂X is either a component of ∂S or an essential, non-peripheral simple closed curve in S (essential means not homotopic to a point or a puncture, and non-peripheral means not homotopic to a component of ∂S). We further assume that no two connected components of S are homotopic to each other (in other words, if X has an annular component, then no other component is a parallel annulus). We also let $\bar{X}$ denote the surface obtained from X by collapsing each component of the boundary to a marked point (we may alternatively regard the marked points as punctures).

For a subgroup G of $\mathrm{Mod}(S)$, let $\mathrm{Stab}_G(X)$ be the stabilizer in G of the homotopy class of X. There is a well-defined map $\mathrm{Stab}_G(X) \to \mathrm{Mod}(\bar{X})$. We denote the image of $f \in \mathrm{Stab}_G(X)$ by $\bar{f}$.

A partial pseudo-Anosov mapping class is an $f \in \mathrm{Mod}(S)$ that has a representative supported on a connected subsurface X and whose image $\bar{f}$ in $\mathrm{Mod}(\bar{X})$ is pseudo-Anosov.

It follows from the Birman–Lubotzky–McCarthy theory of canonical reduction systems [BLM83] that the support X of a partial pseudo-Anosov mapping class f is a well-defined homotopy class of subsurfaces of S.

Independence. We say that pseudo-Anosov mapping classes $f_1, f_2 \in \mathrm{Mod}(S)$ are independent if their corresponding sets of fixed points in $\mathrm{PMF}(\bar{S})$ are disjoint (equivalently, not equal). McCarthy [McC85] proved that f_1 and f_2 are independent in this sense if and only if no non-trivial power of $\bar{f}_1$ is equal to a power of $\bar{f}_2$ (if $\partial S = \emptyset$ then $\bar{S} = S$ and $\bar{f}_i = f_i$).

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Suppose that G is a subgroup of $\mathrm{Mod}(S)$. We further say that $f_1, f_2 \in G$ are G-independent if every conjugate of f_1 by an element of G is independent of f_2 (equivalently, the fixed sets in $\mathrm{PMF}(\bar{S})$ lie in different G-orbits).

The definitions of independence and G-independence carry over to the case where f_1 and f_2 are partial pseudo-Anosov elements with the same support X. Specifically, they are independent if the corresponding maps $\bar{f}_1$ and $\bar{f}_2$ are independent, and they are G-independent if the conjugates by $\mathrm{Stab}_G(X)$ are independent.

We may further extend the definition to the case where f_1 and f_2 are arbitrary partial pseudo-Anosov elements. Specifically, we say that f_1 and f_2 are G-independent if either (1) their supports do not lie in the same G-orbit or (2) their supports do lie in the same G-orbit and conjugates of f_1 and f_2 with the same support are G-independent as in the previous paragraph.

Even further, let f_1 and f_2 be two mapping classes and assume that each f_i is a pseudo-Anosov mapping class, a partial pseudo-Anosov mapping class, or a power of a Dehn twist. Then we say that f_1 and f_2 are G-independent if either (1) their supports lie in different G-orbits or (2) their supports lie in the same G-orbit, they are not both powers of Dehn twists, and they are G-independent as described previously.

Finally, if $\mathcal{F} \subseteq \mathrm{Mod}(S)$ is an arbitrary collection of pseudo-Anosov mapping classes, partial pseudo-Anosov elements, and powers of Dehn twists, then we say that $\mathcal{F}$ is G-independent if each pair of elements of $\mathcal{F}$ is G-independent as described previously.

Elementary closure and NEC mapping classes. Suppose $\partial S = \emptyset$. Again let G be a subgroup of $\mathrm{Mod}(S)$, and let $f \in G$ be a pseudo-Anosov element. The elementary closure in G of f, written $\mathrm{EC}_G(f)$, is defined to be the stabilizer in G of the associated pair of fixed points in $\mathrm{PMF}(S)$. For $G = \mathrm{Mod}(S)$ we simply write $\mathrm{EC}(f)$. We say that f is an NEC element (for 'normal in elementary closure') of G if $\langle f \rangle$ is normal in $\mathrm{EC}_G(f)$.

McCarthy [McC85] proved that $\mathrm{EC}(f)$ is an extension of $\mathbb{Z}$ or D_∞ by a finite subgroup of $\mathrm{Mod}(S)$. As the order of a finite subgroup of $\mathrm{Mod}(S)$ is bounded by a function of S, it follows that there is a $d = d(S)$ so that the dth power of any pseudo-Anosov mapping class is NEC in $\mathrm{Mod}(S)$. A pseudo-Anosov element that is NEC in $\mathrm{Mod}(S)$ is also NEC in an arbitrary subgroup G, and so there is a $d = d(S)$ so that the dth power of any pseudo-Anosov element of any subgroup G is an NEC element of G.

For the case where $\partial S \neq \emptyset$, we say that a pseudo-Anosov $f \in \mathrm{Mod}(S)$ is NEC if $\bar{f} \in \mathrm{Mod}(\bar{S})$ is.

Finally, we may again carry over the definitions to the partial pseudo-Anosov case. Let $f \in \mathrm{Mod}(S)$ be a partial pseudo-Anosov element with support X. Consider the map $\mathrm{Stab}_G(X) \to \mathrm{Mod}(\bar{X})$. Denote the image of $\mathrm{Stab}_G(X)$ by $\bar{G}$ and the image of f by $\bar{f}$. We say that f is an NEC element of G if $\bar{f}$ is an NEC element of $\bar{G}$.

A power of a Dehn twist is, by definition, an NEC element of $\mathrm{Mod}(S)$.

In this paper, we give two further sufficient conditions for a pseudo-Anosov mapping class to be an NEC element of $\mathrm{Mod}(S)$. See the discussion after Proposition 1.4 as well as Lemma 9.3.

Orbit-overlapping subsurfaces. . . . Let X and Y be two homotopy classes of connected subsurfaces of S. We say that X and Y overlap if the boundary of X has essential intersection with Y, and vice versa (meaning that the boundary of every representative of X intersects every representative of Y, and vice versa).

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

If the boundaries of (representatives of) X and Y have essential intersection, then X and Y overlap. However, the converse is not true: consider, for example, the case where X and Y are the complements of disjoint, homotopically distinct non-separating annuli in S.

We may extend the definition of overlapping to the case of subsurfaces that are not necessarily connected. We say that two arbitrary subsurfaces X and Y overlap if each component of X overlaps with each component of Y.

Let G be a subgroup of $\mathrm{Mod}(S)$. We say that a homotopy class of subsurfaces X is G-overlapping if for each $h \in G$ we either have that hX is equal to X or it overlaps with X. We say that a family $\mathcal{X}$ of homotopy classes of subsurfaces of S is G-overlapping if each $X \in \mathcal{X}$ is G-overlapping and for each distinct pair $X, Y \in \mathcal{X}$ any elements of the G-orbits of X and Y overlap. In particular, we must have that X and Y lie in different G-orbits.

One example of a $\mathrm{Mod}(S_g)$-overlapping subsurface of S_g is the complement of a non-separating annulus. Another example of a $\mathrm{Mod}(S_g)$-overlapping subsurface of S_g is an annular neighborhood of a separating curve of genus $g/2$, that is, a separating curve that divides S_g into two surfaces of genus $g/2$ (we may also say that the curve itself is $\mathrm{Mod}(S_g)$-overlapping). Of course, if $G_1 \leqslant G_2$, then a G_2-overlapping family of subsurfaces is G_1-overlapping. Thus, the given examples are G-overlapping for each subgroup $G \leqslant \mathrm{Mod}(S_g)$.

In general, the G-overlapping condition is quite restrictive. For instance, Corollary 2.2 in § 2 states that if $G = \mathrm{Mod}(S)$ and X is a G-overlapping subsurface, then X has at most one annular component and at most two non-annular components. Proposition 2.1 gives further restrictions.

In their work on the large-scale geometry of big mapping class groups, Mann–Rafi [MR20] define the notion of a non-displaceable subsurface of a surface, which is an analogue of our G-overlapping condition in the case that $G = \mathrm{Mod}(S)$. For a connected subsurface, their notion is the same as our notion of $\mathrm{Mod}(S)$-overlapping, but for disconnected subsurfaces it is less restrictive.

...and the associated mapping classes. As previously, let G be a subgroup of $\mathrm{Mod}(S)$. Let $\mathcal{X}$ be a G-overlapping family of homotopy classes of subsurfaces of S. We say that a family of mapping classes $\mathcal{F} \subseteq G$ is carried by $\mathcal{X}$ if there is a function

$$\sigma\colon\ \mathcal{F} \to \mathcal{X}$$

so that the following conditions hold:

(i) each $f \in \mathcal{F}$ is a pseudo-Anosov mapping class, partial pseudo-Anosov mapping class, or a power of a Dehn twist, and the support of f is (the homotopy class of) a component of (a representative of) $\sigma(f)$; and
(ii) for any component X of an element of $\mathcal{X}$, there is an $f \in \mathcal{F}$ and $g \in G$ so that the support of gfg^{-1} is X.

In the case where no two components of elements of $\mathcal{X}$ lie in the same G-orbit, we may replace condition (2) with the simpler condition that for any component X of an element of $\mathcal{X}$, there is an $f \in \mathcal{F}$ whose support is X. Condition (2) is required for the situation where an element of $\mathcal{X}$ has two components X_1 and X_2 in the same G-orbit; say $gX_1 = X_2$. In this case, if the family $\mathcal{F}$ had both an element f_1 with support X_1 and the element $f_2 = gf_1g^{-1}$ with support X_2 (as per the simpler condition), then $\mathcal{F}$ would not be G-independent.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Statement of the main theorem. In order to state our main theorem, we need further notation. First, for a subset A of a group G we denote by $A^{(n)}$ the set $\{a^n \mid a \in A\}$. Also, if G is a group and $A \subseteq G$ we denote by $\langle\langle A \rangle\rangle_G$ the normal closure of A in G.

Suppose that G is a subgroup of $\mathrm{Mod}(S)$, that $\mathcal{X}$ is a G-overlapping family of homotopy classes of subsurfaces of S, and that $\mathcal{F} \subseteq G$ is a family of mapping classes carried by $\mathcal{X}$. Also suppose that $Y = hX$ for some $X \in \mathcal{X}$, $h \in G$. For $n > 0$ we define

$$R_Y = h(\langle\langle(\sigma^{-1}(X))^{(n)}\rangle\rangle_{\mathrm{Stab}_G(X)})h^{-1}.$$

This subgroup of G is well-defined, independent of h. For $X \in \mathcal{X}$, the group R_X is the normal closure in $\mathrm{Stab}_G(X)$ of the nth powers of the elements of $\mathcal{F}$ that are supported in X. For $Y = hX$, we have R_Y is the conjugate of R_X by h. We emphasize that the notation R_Y does not reflect the dependence on $\mathcal{F}$ and n.

Finally, we denote by a_Y and b_Y the numbers of annular and non-annular components of a subsurface $Y \subseteq S$ (or a homotopy class).

Theorem 1.1. *Let $S = S_{g,p}^b$, let G be a subgroup of $\mathrm{Mod}(S)$, let $\mathcal{X}$ be a G-overlapping family of subsurfaces of S, and let $\mathcal{F}$ be a finite, G-independent family of mapping classes that are carried by $\mathcal{X}$ and are NEC in G. There is an $N > 0$ with the following properties:*

(i) *for each $n \geq N$ and any set of orbit representatives $\mathcal{Y}$ for the action of $\langle\langle \mathcal{F}^{(n)} \rangle\rangle_G$ on $G \cdot \mathcal{X}$, we have*

$$\langle\langle \mathcal{F}^{(n)} \rangle\rangle_G \cong \mathop{\ast}_{Y \in \mathcal{Y}} R_Y;$$

(ii) *further, for each $Y \in \mathcal{Y}$ we have*

$$R_Y \cong F_\infty^{b_Y} \times \mathbb{Z}^{a_Y}.$$

As every $f \in \mathrm{Mod}(S)$ has a power p that is NEC in G, and because a power of an NEC mapping class is NEC, we may remove the NEC hypothesis from Theorem 1.1 at the expense of replacing 'for all $n > N$' with 'there exists an N so that for all multiples n of N'; specifically, N is chosen so that each element of $\mathcal{F}^{(N)}$ is NEC in G. In other words, if we remove the NEC hypothesis, the conclusion holds for some specific N (and its multiples) instead of all sufficiently large n.

It follows from the proof that there exists a set of orbit representatives $\mathcal{Y}$ so that the isomorphism in the first statement of Theorem 1.1 is given by inclusion.

Connection to the deep relations question of Ivanov. Let $\mathcal{T}_k(S)$ denote the normal subgroup of $\mathrm{Mod}(S)$ generated by all kth powers of Dehn twists. Ivanov asked [Iva06, § 12] if $\mathcal{T}_k(S)$ has a presentation where the generators are all kth powers of Dehn twists and the relations are the obvious ones, namely the relations

$$T_d^k T_c^k T_d^{-k} = T_{T_d^k(c)}^k.$$

This question was recently answered in the affirmative by Dahmani [Dah18]. In the proof, Dahmani applies a version of the Dahmani–Guirardel–Osin machinery.

One interpretation of our Theorem 1.1 is that there is a presentation for the group $\langle\langle \mathcal{F}^{(n)} \rangle\rangle_G$ where the set of generators is the union of all G-conjugates of $\mathcal{F}^{(n)}$, and where the relations are the obvious ones. Specifically, the generators are of the form hfh^{-1} with $f \in \mathcal{F}^{(n)}$ and $h \in G$.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Denoting each such generator hfh^{-1} as f_h, the relations are all equalities of the form

$$g_k f_h g_k^{-1} = f_{g_k h}.$$

These relations are indeed obvious, as can be seen by expanding them out. Once we have the presentation of the normal closure given by Theorem 1.1, we can obtain the presentation here by repeatedly performing the Tietze transformation of adding a new generator and writing it in terms of the old generators. In this sense, our result is analogous to Dahmani's.

1.3 First examples

We list here some immediate consequences of Theorem 1.1, which exhibit some of the variety of applications. In what follows, $\mathcal{I}(S_g)$ is the Torelli subgroup of $\mathrm{Mod}(S_g)$, which is defined to be the kernel of the action of $\mathrm{Mod}(S_g)$ on $H_1(S_g;\mathbb{Z})$. In addition, we identify the pure braid group on n strands with the pure mapping class group of a disk D_n with n marked points in the interior; see [FM12, Chapter 9].

(i) Taking $\mathcal{X} = \{S\}$ and $G = \mathrm{Mod}(S)$ we obtain the result of Dahmani–Guirardel–Osin that the normal closure of a suitable power of a pseudo-Anosov mapping class is isomorphic to F_∞.

(ii) Taking $\mathcal{X} = \{X\}$ where X is connected and $\mathrm{Mod}(S)$-overlapping and taking $G = \mathrm{Mod}(S)$, we obtain that the normal closure of a suitable power of a partial pseudo-Anosov mapping class with support X is isomorphic to F_∞.

(iii) Taking $\mathcal{X} = \{A\}$ where A is an orbit-overlapping annulus in S, for example an annulus dividing S into two homeomorphic subsurfaces as in Figure 1, and taking $G = \mathrm{Mod}(S)$, we obtain that the normal closure of a suitable power of a Dehn twist about A is isomorphic to F_∞.

(iv) Taking $\mathcal{X} = \{A \cup B\}$ where A is as in example (iii) and B is the complement of an open neighborhood of A (see Figure 1), taking f_A to be a suitable power of a Dehn twist about A, taking f_B to be a suitable power of a partial pseudo-Anosov mapping class supported on one component of B, and taking $G = \mathrm{Mod}(S)$, we obtain that the normal closure of $\{f_A, f_B\}$ in $\mathrm{Mod}(S)$ is isomorphic to

$$\mathop{\ast}_{\infty} \big(F_\infty \times F_\infty \times \mathbb{Z}\big).$$

(v) Taking $\mathcal{X} = \{A\}$ where A is any separating annulus in S_g and taking $G = \mathcal{I}(S_g)$, we obtain that the normal closure in $\mathcal{I}(S_g)$ of a suitable power of a Dehn twist about A is isomorphic to F_∞. Similarly the normal closure in $\mathcal{I}(S_g)$ of a suitable power of any partial pseudo-Anosov element of $\mathcal{I}(S_g)$ is isomorphic to F_∞.

(vi) Taking $\mathcal{X} = \{A\}$ to be any subsurface of D_n and taking G to be the pure braid group, and taking f to be any partial pseudo-Anosov element or Dehn twist supported on A, we obtain that the normal closure of a suitable power of f is isomorphic to F_∞.

(vii) Taking $\mathcal{X} = \{A \cup B \cup C\}$ where A is an annulus in D_n surrounding more than two marked points but fewer than $n-1$ marked points, taking B and C to be the complementary regions to open neighborhood of A, and taking G to be the pure braid group, and taking f_A, f_B, and f_C to be partial pseudo-Anosov elements and Dehn twists with supports equal to A, B, and C, we obtain that the normal closure of suitable powers of f_A, f_B, and f_C is

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

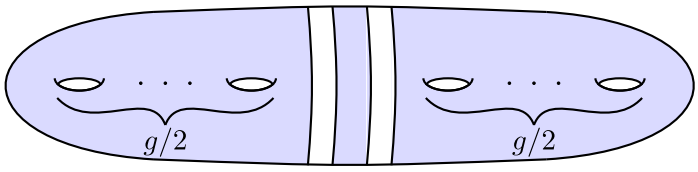

Figure 1. An orbit overlapping subsurface in S_g.

isomorphic to

$$\ast_{\infty} (F_\infty \times F_\infty \times \mathbb{Z}).$$

(viii) Taking $\mathcal{X} = \{A_1 \cup \cdots \cup A_n\}$ where the A_i are pairwise disjoint and each A_i is the support of a partial pseudo-Anosov mapping class f_i and taking G to be the subgroup of $\mathrm{Mod}(S)$ consisting of elements that preserve each A_i, we obtain that the normal closure in G of suitable powers of the f_i is isomorphic to the direct product of n copies of F_∞.

1.4 Application: non-congruence normal subgroups

We give here an application of Theorem 1.1 to the general theory of normal subgroups of the mapping class group. The following theorem answers a question asked in (the first version of) a paper by Lanier and the third author [LM18].

Theorem 1.2. *Let $g \geq 2$. There exists a proper normal subgroup of $\mathrm{Mod}(S_g)$ that is not contained in any proper level m congruence subgroup of $\mathrm{Mod}(S_g)$.*

To prove Theorem 1.2, we first choose independent NEC pseudo-Anosov mapping classes f_1 and f_2, each of whose actions on $H_1(S_g;\mathbb{Z})$ is equal to that of a Dehn twist about some non-separating curve. Then for distinct large primes p_1 and p_2, the normal closure N of $\{f_1^{p_1}, f_2^{p_2}\}$ is the desired subgroup. Indeed, Theorem 1.1 implies N is free (hence, proper) and it is evident that N is not contained in any level m congruence subgroup. The main work is to prove that f_1 and f_2 exist. In the process, we make f_1 and f_2 (but not p_1 and p_2) explicit.

After learning about our work, Ashot Minasyan pointed out to us that Theorem 1.2 can also be derived from earlier work of Hull [Hul16]. We explain the details of this argument in § 9.

1.5 Which RAAGs?

Putting the two statements of Theorem 1.1 together and applying the restrictions on a_Y and b_Y from Corollary 2.2 (mentioned previously), we see that when $G = \mathrm{Mod}(S)$ the group $\langle\langle \mathcal{F}^{(n)} \rangle\rangle_G$ is isomorphic to a free product of the groups

$$F_\infty, \quad \ast_{\infty} (F_\infty \times \mathbb{Z}), \quad \ast_{\infty} (F_\infty \times F_\infty), \quad \text{and} \quad \ast_{\infty} (F_\infty \times F_\infty \times \mathbb{Z});$$

We prove in § 2 that all possibilities occur, that is, any group that is a free product of groups

$$\ast_{\infty} F_\infty^{b_i} \times \mathbb{Z}^{a_i},$$

where $a_i \in \{0,1\}$ and $b_i \in \{0,1,2\}$ for all i, is isomorphic to some $\langle\langle \mathcal{F}^{(n)} \rangle\rangle_{\mathrm{Mod}(S)}$; see Theorem 2.3. In particular, this gives the first examples of normal, non-free, right-angled Artin subgroups of the mapping class group.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Theorem 1.1 applies to mapping classes whose supports are 'sufficiently large'. On the other hand, a result of Brendle and the third author roughly states that the normal closure of any mapping class with 'sufficiently small' support is not isomorphic to a RAAG [BM19, Corollary 1.4]. This leads us to the following conjecture.

Conjecture 1.3. Let $S = S_{g,p}$. If N is a non-trivial normal subgroup of $\mathrm{Mod}(S)$ and N is isomorphic to a RAAG then N is isomorphic to one of the right-angled Artin subgroups of $\mathrm{Mod}(S)$ afforded by our construction. In particular, N is isomorphic to a free product of the groups

$$F_\infty, \quad \ast_\infty (F_\infty \times \mathbb{Z}), \quad \ast_\infty (F_\infty \times F_\infty), \quad \text{and} \quad \ast_\infty (F_\infty \times F_\infty \times \mathbb{Z}).$$

For specific values of g and n the conjecture may be sharpened. For instance, when $S = S_g$ the conjecture says that the RAAG is isomorphic to a free product of the groups

$$F_\infty, \quad \ast_\infty (F_\infty \times F_\infty), \quad \text{and} \quad \ast_\infty (F_\infty \times F_\infty \times \mathbb{Z}),$$

because the group $\ast (F_\infty \times \mathbb{Z})$ does not arise from our construction in this case (this is a consequence of Proposition 2.1); cf. Theorem 2.4. Similarly, when at least one of g or n is odd, the group $\ast (F_\infty \times F_\infty \times \mathbb{Z})$ does not appear, etc.

Other normal free groups.... Farb explicitly asked whether or not a pseudo-Anosov mapping class has a power with free normal closure [Far06, Question 2.9]. Before the work of Dahmani–Guirardel–Osin, the only known examples of free, normal subgroups of the mapping class group were those due to Whittlesey [Whi00]. She proved that the Brunnian subgroup of $\mathrm{Mod}(S_{0,n})$ is free. As a consequence, she proved that there is a corresponding free, normal, all pseudo-Anosov subgroup of $\mathrm{Mod}(S_2)$.

...and non-normal RAAGs. Long before the work of Dahmani–Guirardel–Osin, it was proven that if two curves have geometric intersection number greater than 1 then the corresponding Dehn twists generate a free group of rank two. This result was proved by Ishida [Ish96] and Hamidi–Tehrani [HT02], and also appears in Handel's notes from Thurston's course on mapping class groups in Princeton from 1975. Ivanov [Iva92, Corollary 8.4] and McCarthy [McC85] proved that high powers of independent pseudo-Anosov mapping classes generate a free group of rank two. Since then there have been many different constructions of non-normal free subgroups of the mapping class group with various properties; see, for instance, the work of Fujiwara [Fuj15] and of the second author [Man10, Man13].

Beyond free groups, there are a number of papers devoted to the problem of finding right-angled Artin subgroups of mapping class groups. Constructions of such subgroups are given by Leininger and the first two authors of this paper [CLM12], Crisp–Farb [CF], Crisp–Paris [CP01], Crisp–Wiest [CW07], Koberda [Kob12], Lönne [Lön10], Runnels [Run20], and Seo [Seo19]. Each of these works has its own points of emphasis, but one over-arching theme is that every finitely generated RAAG is isomorphic to a subgroup of some mapping class group. In all of these works, the resulting subgroup is not normal.

1.6 Which powers?

As mentioned previously, the special case of Theorem 1.1 where each element of $\mathcal{F}$ is pseudo-Anosov is due to Dahmani–Guirardel–Osin. The most obvious distinctions between their work

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

and ours are that Theorem 1.1 applies to certain types of reducible mapping classes, and also that our normal closures are not always free groups. Beyond this, we help clarify the situation as to which powers of a pseudo-Anosov mapping class do, and do not, have free normal closure. We now discuss these two points in more detail.

Which powers do have free normal closure?. In a forthcoming paper we apply the techniques of this paper to show the following.

(i) If $f \in \mathrm{Mod}(S_{g,p})$ is an NEC pseudo-Anosov element, then $\langle\langle f^N \rangle\rangle$ is free for

$$N \geq \exp\left(\exp\left(10^6 \cdot \delta_{g,p}^2 \cdot (3g-3+p)^2\right)\right).$$

(ii) If c is an orbit overlapping curve and $p = 0$, then $\langle\langle T_c^N \rangle\rangle$ is free for $N \geq 53{,}489$.

Here, $\delta_{g,p}$ is any hyperbolicity constant for the curve graph of $S_{g,p}$; it has been shown that $\delta_{g,p}$ can be taken to be 17, independently of $S_{g,p}$ (see [HPW15]).

It is an interesting problem to sharpen these values of N. For the case of an orbit overlapping curve, we can see that $N = 1$ does not suffice to ensure that $\langle\langle T_c^N \rangle\rangle$ is free. For example, when c is a curve of genus $g/2$ in S_g, the lantern relation can be used to exhibit the non-freeness (if a, b, and c are curves of genus $g/2$ that lie in a four-holed sphere L in S_g, then by the lantern relation we have $T_a T_b = T_c^{-1} M$ where M is the multitwist about the boundary of L, from which we see that $T_a T_b$ and T_c commute but do not have a common power). On the other hand, we do not know whether or not $\langle\langle T_c^2 \rangle\rangle$ is free. Similarly, $N = 1$ does not suffice in the pseudo-Anosov case: as explained in Proposition 1.4, every pseudo-Anosov element of the Torelli group $\mathcal{I}(S_g)$ is NEC, and Lanier and the third author proved [LM20, Theorem 1.3] that there is a pseudo-Anosov element of $\mathcal{I}(S_g)$ whose normal closure in $\mathrm{Mod}(S_g)$ is equal to $\mathcal{I}(S_g)$.

Which powers do not have free normal closure?. It follows from the work of Dahmani–Guirardel–Osin, specifically Corollary 6.36 and Theorem 6.8 in their paper [DGO17], that sufficiently large powers of NEC pseudo-Anosov mapping classes have free normal closure (this is also a special case of Theorem 1.1). Many results about pseudo-Anosov mapping classes hold for sufficiently large powers, for example, the result of Ivanov and McCarthy mentioned previously. Thus, one may be tempted to think that all sufficiently large powers of a pseudo-Anosov mapping class have free normal closure. The following proposition shows that this is not the case.

Proposition 1.4. *Let $f \in \mathrm{Mod}(S)$ be a pseudo-Anosov mapping class that is not NEC. Then for arbitrarily large n, the group $\langle\langle f^n \rangle\rangle$ contains a non-trivial periodic element; in particular, $\langle\langle f^n \rangle\rangle$ is not free.*

Indeed, any $h \in \mathrm{EC}(f)$ conjugates f to $hfh^{-1} = f^{\pm 1} r$ for some r in the finite group of symmetries of the singular Euclidean structure on S associated to f, as these are precisely the pseudo-Anosov elements in $\mathrm{EC}(f)$ with the same stretch factor as f (this follows from McCarthy's description of $\mathrm{EC}(f)$ in his unpublished paper [McC85]). If f is not NEC, then some h gives a non-trivial $r = f^{\mp 1} h f h^{-1}$ in $\langle\langle f \rangle\rangle$. Moreover, if f is not NEC, then arbitrarily large powers of f are not NEC, hence the proposition.

It follows immediately from Proposition 1.4 that if f lies in a normal subgroup of $\mathrm{Mod}(S)$ that is torsion free, then f is NEC. As an example, if f lies in the level m congruence subgroup

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

of $\mathrm{Mod}(S_g)$ with $m \geq 3$, then f is NEC in $\mathrm{Mod}(S_g)$. In particular, every pseudo-Anosov element of the Torelli group $\mathcal{I}(S_g)$ is NEC.

Proposition 1.4 shows that the NEC condition is necessary in order for a pseudo-Anosov mapping class to have free normal closure. However, it is not sufficient. Indeed, as mentioned previously, Lanier and the third author gave an example of a NEC pseudo-Anosov mapping class with normal closure $\mathcal{I}(S_g)$. On the other hand, the following question seems to be open: is the normal closure of an NEC pseudo-Anosov mapping class torsion free?

In the case of a closed surface, we can strengthen the conclusion of Proposition 1.4: not only does $\langle\langle f^n \rangle\rangle$ fail to be free, but it fails to be abstractly commensurable to any RAAG. Indeed, it follows from work of Lanier and the third author [LM20, Theorem 1.1] that the normal closure of any non-trivial periodic element of $\mathrm{Mod}(S_g)$, hence the normal closure of f^n, contains $\mathcal{I}(S_g)$. In addition, Brendle and the third author proved that the normal closures of any subgroup of $\mathrm{Mod}(S_g)$ containing $\mathcal{I}(S_g)$ is not abstractly commensurable with any RAAG [BM19, Corollary 1.4].

The failure of $\langle\langle f^n \rangle\rangle$ to be free (or abstractly commensurable to any RAAG) is underscored by the following result, which we prove in § 9.

Theorem 1.5. *For each $g \geq 3$ there exists a pseudo-Anosov $f \in \mathrm{Mod}(S_g)$ so that if n is odd then*

$$\langle\langle f^n \rangle\rangle = \mathrm{Mod}(S_g),$$

and if n is even and non-zero, then

$$\langle\langle f^n \rangle\rangle \cong F_\infty.$$

The mapping classes that appear in the proof of Theorem 1.5 are based on those used by Lanier and the third author to show that there are pseudo-Anosov elements of $\mathrm{Mod}(S_g)$ with the property that all of their odd powers are normal generators (in particular, there are normal generators for $\mathrm{Mod}(S_g)$ with arbitrarily large translation lengths on the curve complex) [LM20, Theorem 1.4].

1.7 Windmills and spinning families

Our next goal is to explain the main technical theorem used to prove Theorem 1.1, namely, Theorem 1.6. This theorem concerns the theory of group actions on projection complexes. Briefly, a projection complex is a graph Γ that comes equipped with a collection of functions

$$d_v \colon V \setminus \{v\} \times V \setminus \{v\} \to \mathbb{R}_{\geq 0},$$

where V is the set of vertices of Γ and $v \in V$. Projection complexes were defined by Bestvina–Bromberg–Fujiwara; see § 3 for the full definition, along with examples and motivation. Our definition is a mild modification of the original definition of Bestvina–Bromberg–Fujiwara in that we require the projection complex to satisfy additional properties (such as the bounded geodesic image property).

After stating Theorem 1.6, we explain forthcoming work on the geometry of the corresponding quotient complexes, and then we state Theorem 1.7. The latter is the simplest type of application of Theorem 1.6, where the output is a free group. In § 1.8 we explain how Theorems 1.6 and 1.7 are pieced together to prove Theorem 1.1.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Windmills in projection complexes. Let $\mathcal{P}$ be a projection complex, and let G be a group that acts on $\mathcal{P}$. Further, for each vertex v of $\mathcal{P}$, let R_v be a subgroup of the stabilizer of v in G. Let $L > 0$. We say that the family of subgroups $\{R_v\}$ is an *equivariant L-spinning family* of subgroups of G if it satisfies the following two conditions.

- *Equivariance:* if g lies in G and v is a vertex of $\mathcal{P}$, then
$$gR_vg^{-1} = R_{gv}.$$
- *Spinning condition:* for any distinct vertices v and w of $\mathcal{P}$ and any non-trivial $h \in R_v$ we have
$$d_v(w, hw) \geq L.$$

The equivariance condition implies that for each vertex v, the subgroup R_v is normal in $\mathrm{Stab}(v)$, and that the subgroup $\langle R_v \rangle$ of G generated by the R_v is normal in G. Furthermore, if we choose orbit representatives $\{v_i\}$ for the action of G on the vertices of $\mathcal{P}$, then $\langle R_v \rangle$ is the normal closure of the set $\{R_{v_i}\}$.

Theorem 1.6. *Let $\mathcal{P}$ be a projection complex and let G be a group acting on $\mathcal{P}$. There exists a constant $L(\mathcal{P})$ with the following property. If $L \geq L(\mathcal{P})$, if $\{R_v\}$ is an equivariant L-spinning family of subgroups of G, and if $\mathcal{O}$ is any set of orbit representatives for the action of $\langle R_v \rangle$ on the set of vertices of $\mathcal{P}$, then $\langle R_v \rangle$ is isomorphic to the free product*
$$\ast_{v \in \mathcal{O}} R_v.$$

In the statement of Theorem 1.6, we emphasize that $\langle R_v \rangle$ is the group generated by all of the R_v, not just one R_v. Also, it follows from the proof that there exists a set of orbit representatives $\mathcal{O}$ so that the isomorphism in Theorem 1.6 is given by inclusion.

To prove Theorem 1.6 we introduce the notion of a windmill in a projection complex; see § 4. A windmill is a subgraph of $\mathcal{P}$ that serves as a proxy for the Bass–Serre tree for the desired free product decomposition. Our definition of a windmill is derived from the theory of windmills for group actions on hyperbolic spaces, due to Dahmani–Guirardel–Osin, which, in turn, has its origins in the work of Gromov [Gro01, § 26].

As we explain in § 3, every simplicial tree can be regarded as a projection complex (and, conversely, a projection complex is a quasi-tree [BBF15, Theorem 3.16]). In § 5 we give the proof of Theorem 1.6 in the special case where $\mathcal{P}$ is a tree, as a warmup for the general case.

After the first version of this paper appeared, Bestvina–Dickmann–Domat–Kwak–Patel–Stark gave a new proof of Theorem 1.6. Their proof produces an $\langle R_v \rangle$-invariant tree in $\mathcal{P}$ that is the Bass–Serre tree for the given free product decomposition. In their paper [BDDKPS20], they also explain how to derive a version of the main result of the paper by Dahmani–Guirardel–Osin [DGO17, Theorem 5.3a] from Theorem 1.6.

Hyperbolicity of the quotient. As previously, $\mathcal{T}_k(S)$ denotes the normal subgroup of $\mathrm{Mod}(S)$ generated by kth powers of Dehn twists. Building on the aforementioned work of Dahmani, it was recently shown by Dahmani–Hagen–Sisto that for suitable k the quotient group $\mathrm{Mod}(S)/\mathcal{T}_k(S)$ is acylindrically hyperbolic [DHS21] (in particular, the group $\mathcal{T}_k(S)$ has infinite index in $\mathrm{Mod}(S)$). One of the major steps in the proof is to show that the quotient of the curve complex by $\mathcal{T}_k(S)$ is hyperbolic.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

In analogy with this result, the first two authors prove in a separate paper that under certain hypotheses, the quotient complex $\mathcal{P}/\langle R_v \rangle$ arising in Theorem 1.6 is hyperbolic and that the quotient group $G/\langle R_v \rangle$ is acylindrically hyperbolic [CM]. The argument follows along the lines of the work of Dahmani–Hagen–Sisto [DHS21].

Free groups from windmills. Our next result, Theorem 1.7, is the simplest type of application of Theorem 1.6, in that the R_v subgroups are all isomorphic to $\mathbb{Z}$ (and so the output is F_∞). Theorem 1.7 also serves as a sort of base case for Theorem 1.1; in the outline of the paper given in the following, we explain this in more depth.

For the statement we require several definitions. First, when we say that a space is hyperbolic, we mean that it is a metric space that is δ-hyperbolic in the sense of Gromov. Fix some hyperbolic space X and suppose a group G acts on X by isometries. Certain types of elements of G are called WPD elements (for 'weakly properly discontinuous'). The idea is that f is WPD if its translation length is positive and if long segments of a quasi-axis for f have a finite coarse stabilizer; see § 3.2 for the details. One important example for our work is where $G = \mathrm{Mod}(S)$ and $X = \mathcal{C}(S)$; in this case, the WPD elements are exactly the pseudo-Anosov elements of $\mathrm{Mod}(S)$.

Each WPD element $f \in G$ has two fixed points in ∂X, and so the elementary closure can be defined for f in the same way as for pseudo-Anosov elements of $\mathrm{Mod}(S)$: it is the stabilizer of this pair of points in ∂X. We say that f is NEC if it is normal in its elementary closure. In addition, we say that f_1 and f_2 are independent if $\mathrm{EC}(f_1) \cap \mathrm{EC}(f_2)$ is finite and they are normally independent if every conjugate of f_1 is independent of f_2. Finally, a family of WPD elements is normally independent if they are pairwise normally independent.

Theorem 1.7. *Suppose G acts on a hyperbolic space X and $\{f_1, \ldots, f_m\} \subseteq G$ is a normally independent collection of NEC WPD elements of G. For any $t \geq 0$ there is a constant N such that for $n \geq N$ the group*

$$\langle\langle f_1^n, \ldots, f_m^n \rangle\rangle_G$$

has the following properties:

(i) *it is isomorphic to F_∞ with a free basis consisting of conjugates of the f_i^n; and*
(ii) *the translation length of each non-trivial element is at least t.*

As before, every WPD element has a power that is NEC and so Theorem 1.7 may be applied to any normally independent collection of WPD elements after replacing each by a power.

In addition to the mapping class group, another application of Theorem 1.7 is to the outer automorphism group of a finitely generated free group F_n. In this case the space X is the free factor complex of F_n, and the collection $\{f_1, \ldots, f_m\} \subseteq \mathrm{Out}(F_n)$ is any collection of normally independent fully irreducible outer automorphisms. See the paper by Bestvina–Feighn [BF14] for the relevant definitions.

1.8 Outline of the paper

We begin in § 2 by describing the basic properties of orbit-overlapping families of subsurfaces. In particular, in § 2.1 we classify all orbit-overlapping families for an arbitrary normal subgroup of the mapping class group. We also show in Theorem 2.3 that all of the RAAGs discussed before do indeed appear as normal subgroups of some mapping class group. Putting these two results together, we obtain a complete classification of right-angled Artin subgroups than can arise from

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

our construction. We also show that the same RAAGs appear as normal subgroups of pure braids groups and Torelli groups.

In § 3 we give the definition of a projection complex. After explaining the basic examples, we present a new example based on disconnected subsurfaces that is used for the proof of Theorem 1.1. We conclude this section with a construction of Bestvina–Bromberg–Fujiwara and Dahmani–Guirardel–Osin that allows us to build a projection complex from a group action on a hyperbolic space. For the proof of Theorem 1.7 we apply their construction to the case of the $\mathrm{Mod}(S)$ action on the curve graph $\mathcal{C}(S)$; the key fact about this action that is needed for the construction is that the action of a pseudo-Anosov mapping class on $\mathcal{C}(S)$ is WPD.

Section 4 provides the architecture for the proof of Theorem 1.6, and gives two successive reductions of Theorem 1.6. We begin by defining the windmill associated to the action of a group G on a projection complex $\mathcal{P}$. A windmill is the direct limit of an increasing union of certain subgraphs W_i. We then define an increasing sequence of free products F_i, each equipped with a homomorphism to G and an action on W_i. Proposition 4.1 states that Theorem 1.6 holds if each homomorphism $F_i \to G$ is injective; this is the first reduction. A natural way to prove the injectivity is to show that the induced action on $\mathcal{P}$ is faithful. In § 4 we introduce the notions of pivot points and waypoints as tools for showing that this action is faithful (Proposition 4.2).

In § 5 we use the machinery developed in § 4 to prove Theorem 1.6 in the special case where the projection complex is a tree. The special case is stated as Theorem 5.1. While this theorem is an immediate consequence of the Bass–Serre theory for group actions on trees, the proof serves as a demonstration of the windmill machinery in action, and also serves as a template for the proof of Theorem 1.6 in the following section.

Section 6 contains the proof of Theorem 1.6, our main technical result about projection complexes. The proof is analogous to the proof of Theorem 5.1, but it is more complicated. The main new difficulty is that the graphs W_i used to define the windmill are not necessarily convex (in a tree any connected subset is convex).

We next turn toward the proof of Theorem 1.1. First, in § 7 we prove Theorem 1.7. This theorem gives the special case of Theorem 1.1 where $\mathcal{X} = \{S\}$; this special case serves as a sort of base case for Theorem 1.1. The idea of the proof of Theorem 1.7 is to use the aforementioned construction of Bestvina–Bromberg–Fujiwara and Dahmani–Guirardel–Osin to translate the given action by WPD elements to an equivariant spinning action on a projection complex, and then to apply our result about actions on projection complexes, Theorem 1.6.

In § 8 we build on Theorems 1.6 and 1.7 in order to prove Theorem 1.1. The first step is to build a projection complex from the given G-overlapping family of subsurfaces $\mathcal{X}$; the vertices correspond to the subsurfaces of S lying in the G-orbits of the elements of $\mathcal{X}$. The L-spinning condition for this action is then ensured by Theorem 1.7 and the orbit overlapping condition. With this in place, we may apply Theorem 1.6 to obtain Theorem 1.1. We note that Theorem 1.6 is applied twice overall, once in the proof of Theorem 1.1 and once in the proof of Theorem 1.7 (which is then used in the proof of Theorem 1.1).

In § 9 we prove two applications of Theorem 1.1, namely, Theorems 1.2 and 1.5. In addition to our theorem, both proofs use the Thurston construction for pseudo-Anosov mapping classes, and the theory of elementary closures of pseudo-Anosov mapping classes developed by McCarthy.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

2. Classification of the RAAGs that arise from Theorem 1.1

We begin in § 2.1 by restricting the RAAGs that could arise from Theorem 1.1 in the case when G is normal in $\mathrm{Mod}(S)$; there are 15 possibilities. In § 2.2 we show that all 15 possibilities do indeed arise for $\mathrm{Mod}(S_{g,p})$ with $g, p \geq 6$ and then we show that certain specific RAAGs arise in the closed case, where $p = 0$. Then in § 2.3 we show that all 15 possibilities also arise in the cases of the pure braid group and the Torelli group.

2.1 Overlapping subsurfaces

Although any subsurface $X \subseteq S$ is G-overlapping when $G = \mathrm{Stab}(X)$, when G is a normal subgroup of $\mathrm{Mod}(S)$ there are topological conditions that must be satisfied by any G-overlapping subsurface. These conditions then impose restrictions on the RAAGs that could arise from Theorem 1.1 when G is normal in $\mathrm{Mod}(S)$.

Proposition 2.1. *Suppose that $G \trianglelefteq \mathrm{Mod}(S)$, that X is a G-overlapping subsurface, that X_1 and X_2 are arbitrary components of X, and that there is a partial pseudo-Anosov element or a power of a Dehn twist in G with support contained in X_1. Then every component of the boundary of X_1 either lies in ∂S or is homotopic to a component of the boundary of X_2.*

Proof. Suppose for the sake of contradiction that some component d of ∂X_1 is not homotopic to a component of the boundary of S or X_2. It follows that there is a curve c that intersects d but not X_2. Let f be a partial pseudo-Anosov mapping class or a Dehn twist with support contained in X_1. By modifying c if needed, we may further assume that c intersects the support of f. For any m and n, the element $g = T_c^{-n} f^m T_c^n$ lies in G as G is normal in $\mathrm{Mod}(S)$. For all choices of m and n, we have that $gX_2 = X_2$, as the support of g is contained in the complement of X_2. We show for large m and n that gX_1 is not equal to X_1. This will mean that gX is not equal to X and that gX does not overlap X, in contrast to the hypothesis that X is G-overlapping.

To show that gX_1 is not equal to X_1 for large enough m and n, we show that $g(d)$ intersects d for large enough m and n. Applying T_c^n to both curves $g(d)$ and d, this is equivalent to the statement that $f^m T_c^n(d)$ intersects $T_c^n(d)$. For large n, $T_c^n(d)$ intersects the support of f (this follows from [FM12, Proposition 3.4]), and then for large m it follows that $f^m(T_c^n(d))$ intersects $T_c^n(d)$ (in the case that f is a power of a Dehn twist use the formula [FM12, Proposition 3.4] again and if f is a partial pseudo-Anosov element use the fact due to Masur–Minsky [MM99, Proposition 4.6] that f has positive translation on the curve complex for the support of f and the fact that the subsurface projection [MM00] of $T_c^n(d)$ to the support of f is well defined). This completes the proof. □

Corollary 2.2. *Suppose that $G \trianglelefteq \mathrm{Mod}(S)$, that X is a G-overlapping subsurface of S, and that there is a partial pseudo-Anosov element or a power of a Dehn twist in G with support contained in X. Then, $0 \leq a_X \leq 1$ and $0 \leq b_X \leq 2$.*

Proof. As a subsurface of S does not have parallel annuli, it follows from Proposition 2.1 that $a_X \leq 1$. Suppose $b_X \geq 2$ and let X_1 and X_2 be two non-annular components of X. It follows from Proposition 2.1 that the complement in S of $X_1 \cup X_2$ is the union of annuli. Thus $b_X = 2$, as desired. □

We say that a subsurface X of S is *compatible* with a subsurface Y if there is a homeomorphism f of S such that $f(X) \subseteq Y$; otherwise, we say that X is *incompatible* with Y. We say

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

that X and Y are *mutually incompatible* if neither is compatible with the other. Similarly, for a subgroup G of $\mathrm{Mod}(S)$ we say that X is *G-compatible* with Y if this f can be chosen so that its homotopy class lies in G; the terms *G-incompatible* and *mutually G-incompatible* are also defined analogously.

We say that a multicurve C is a *dividing set* if $S - C$ consists of two components and $S - C'$ is connected for any proper submulticurve $C' \subset C$. In particular, any separating curve is a dividing set. Given a dividing set C, we let A be a closed regular neighborhood of C and we let L and R be the two components of the complement of an open annular neighborhood of A.

By Corollary 2.2 there are five possibilities for the pair (a_X, b_X) for a G-overlapping subsurface X, where G satisfies the hypotheses of the corollary. Using Proposition 2.1 we can explicitly describe X in each of the five cases, as follows:

- $(0,1)$ X is a connected subsurface that is incompatible with its complementary region;
- $(0,2)$ $X = L \cup R$ for a dividing set C with L and R either homeomorphic or mutually incompatible;
- $(1,0)$ $X = A$ for a separating curve C with L and R either homeomorphic or mutually incompatible;
- $(1,1)$ $X = L \cup A$ for a separating curve C with L and R mutually incompatible;
- $(1,2)$ $X = L \cup A \cup R$ for a separating curve C with L and R either homeomorphic or mutually incompatible.

Let G be a normal subgroup of $\mathrm{Mod}(S)$ that satisfies the hypotheses of Corollary 2.2. By the corollary, the only RAAGs that may result from an application of Theorem 1.1 are of the form

$$\mathop{\ast}_{(a,b)\,\in\, A} \Big(\mathop{\ast}_{\infty} \big((F_\infty)^b \times \mathbb{Z}^a \big) \Big),$$

where A is a subset of $\{0,1\} \times \{0,1,2\}$. As an infinite free product of infinite cyclic groups is isomorphic to an infinite free product of infinitely generated free groups, we may assume that A does not contain, say, $(0,1)$. In summary, there are 15 non-trivial isomorphism types of right-angled Artin subgroups of G that may arise from an application of our Theorem 1.1, and they correspond to the 15 non-empty subsets of $\{(1,0),(1,1),(0,2),(1,2)\}$.

2.2 Mapping class groups

In this section, we prove that in certain mapping class groups, all 15 of the RAAGs from § 2.1 arise as normal subgroups. Then we determine which of these 15 groups arise in the case that the surface is closed.

Theorem 2.3. *Let $g, p \geq 6$. For each of the 15 non-empty subsets A of $\{(1,0),(1,1),(0,2),(1,2)\}$ there is a normal subgroup of $\mathrm{Mod}(S_{g,p})$ isomorphic to*

$$\mathop{\ast}_{(a,b)\,\in\, A} \Big(\mathop{\ast}_{\infty} \big((F_\infty)^b \times \mathbb{Z}^a \big) \Big).$$

Proof. Let C_1, C_2, C_3 and C_4 be the four curves in $S_{g,p}$ as indicated in Figure 2 for the case of $S_{6,6}$; for $g > 6$ and $p > 6$ we may simply add more handles to the region with three handles and more marked points to the region with three marked points.

Let A_i be a closed annular neighborhood of C_i and let L_i and R_i be the components of the complement of an open annular neighborhood of A_i for $i = 1, \ldots, 4$. By construction, any pair of

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

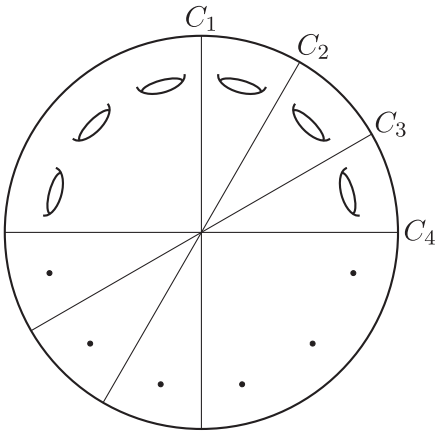

FIGURE 2. The separating curves in Theorem 2.3.

subsurfaces in $\{L_i \mid i = 1, \dots, 4\} \cup \{R_i \mid i = 1, \dots 4\}$ is mutually incompatible with the possible exception of L_1 and R_1 (which are homeomorphic when $g = p = 6$). From this it follows that

$$\mathcal{X} = \{A_1, L_2 \cup A_2, L_3 \cup R_3, L_4 \cup A_4 \cup R_4\}$$

is $\mathrm{Mod}(S_{g,p})$-overlapping.

Applying Theorem 1.1 to all possible subsets of $\mathcal{X}$ (and choosing $\mathcal{F}$ appropriately) we obtain the desired conclusion.

More specifically, each element (a, b) of $\{(1, 0), (1, 1), (0, 2), (1, 2)\}$ corresponds to an element $X_{(a,b)} \in \mathcal{X}$ as follows:

$$\begin{aligned}(1,0) &\rightsquigarrow A_1,\\ (1,1) &\rightsquigarrow L_2 \cup A_2,\\ (0,2) &\rightsquigarrow L_3 \cup R_3,\\ (1,2) &\rightsquigarrow L_4 \cup A_4 \cup R_4.\end{aligned}$$

We choose suitable mapping classes for each component of each such $X_{(a,b)}$. In each case, we may use the standard fact that a surface $S_{g',p'}$ containing an essential curve has pseudo-Anosov elements in its mapping class group; in particular each L_i and R_i that appears in $\mathcal{X}$ is the support of some partial pseudo-Anosov element. □

THEOREM 2.4. *Let $g \geq 2$.*

(i) *If g is even, there are normal subgroups of $\mathrm{Mod}(S_g)$ isomorphic to*

$$F_\infty, \quad \underset{\infty}{\ast} (F_\infty \times F_\infty \times \mathbb{Z}), \quad \underset{\infty}{\ast} (F_\infty \times F_\infty),$$

and also the free product of the third with either of the first two if $g \geq 4$.

(ii) *If g is odd, there are normal subgroups of $\mathrm{Mod}(S_g)$ isomorphic to*

$$F_\infty, \quad \underset{\infty}{\ast} (F_\infty \times F_\infty),$$

and also the free product of these.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

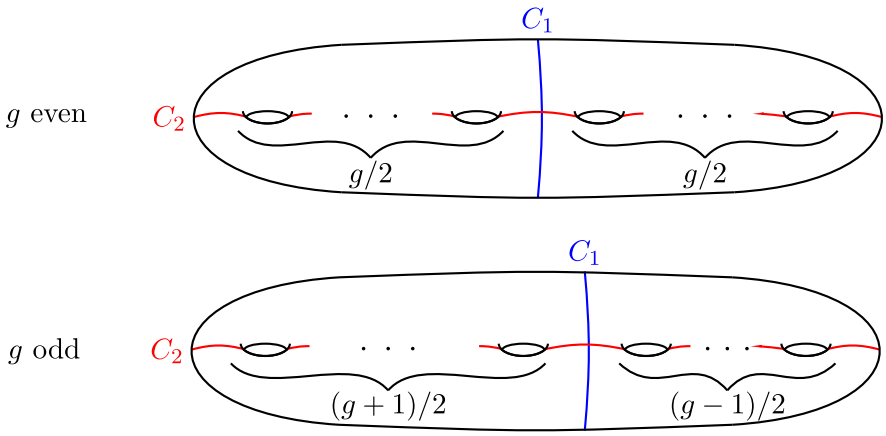

Figure 3. (Colour online) The dividing sets for Theorem 2.4.

Moreover, these are the only normal right-angled Artin subgroups of $\mathrm{Mod}(S_g)$ *that arise from the construction of Theorem 1.1.*

Proof. We begin with the first two statements. For the two cases (g even and g odd), let C_1 and C_2 be the multicurves indicated in the top and bottom of Figure 3, respectively. We define A_i, L_i, and R_i as in the proof of Theorem 2.3.

The desired subgroups are constructed in the same way as in the proof of Theorem 2.3. For instance, when g is even we may construct the free product of $\ast_\infty (F_\infty \times F_\infty \times \mathbb{Z})$ with $\ast_\infty (F_\infty \times F_\infty)$ by taking

$$\mathcal{X} = \{A_1 \cup L_1 \cup R_1, L_2 \cup R_2\},$$

and we may construct the free product of F_∞ with $\ast_\infty (F_\infty \times F_\infty)$ by taking

$$\mathcal{X} = \{A_1, L_2 \cup R_2\}.$$

Similarly, when g is odd we may construct, for example, the free product F_∞ with $\ast_\infty (F_\infty \times F_\infty)$ by taking

$$\mathcal{X} = \{L_1, L_2 \cup R_2\}.$$

The other RAAGs from the statement of the theorem are obtained by removing elements (or components of elements) of the given sets $\mathcal{X}$.

The last statement follows from a case-by-case analysis. For example, the group $\ast_\infty (F_\infty \times \mathbb{Z})$ does not arise because there does not exist a separating curve in S_g with the property that the two complementary regions are mutually incompatible. For the other cases, the reasoning is similar. □

2.3 Pure braid groups and Torelli groups

The idea behind the examples in § 2.2 can be extended to create examples in both the pure braid group PB_n and the Torelli group $\mathcal{I}(S_g)$.

For the pure braid group, we require some setup. A subsurface X of D_n induces a partition of the marked points of D_n: two marked points are in the same subset if and only if there is

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

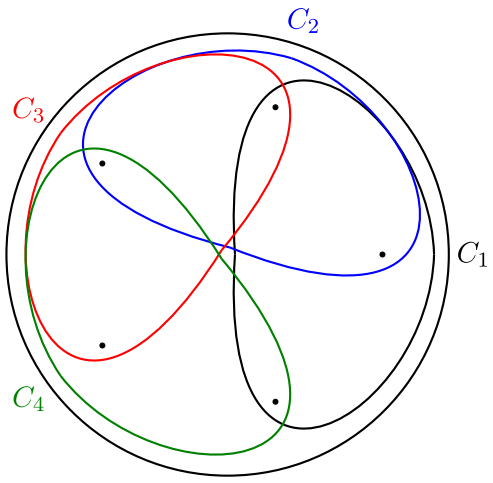

Figure 4. (Colour online) The curves in the proof of Theorem 2.5.

a path between them that does not cross the boundary of X. We say that two partitions of a set are *mutually compatible* if each component of one partition either contains or is contained in a component of the other partition. It is a fact that two subsurfaces of D_n are mutually PB_n-incompatible if the corresponding partitions of the marked points are mutually incompatible.

Theorem 2.5. *Suppose $n \geq 5$. For each of the 15 non-empty subsets A of*

$$\{(1,0),(1,1),(0,2),(1,2)\},$$

there is a normal subgroup of PB_n isomorphic to

$$\underset{(a,b)\,\in\,A}{\ast}\Big(\underset{\infty}{\ast}\big((F_\infty)^b \times \mathbb{Z}^a\big)\Big).$$

Proof. Let C_1, C_2, C_3, and C_4 be the four curves in D_n indicated in Figure 4. The figure shows the case $n = 5$; for $n > 5$ we may simply add marked points in the exterior of all four curves. We define A_i, L_i, and R_i as in the proof of Theorem 2.3.

Consider the following set of subsurfaces of D_n:

$$\{A_1, A_2, A_4\} \cup \{L_2, L_3, L_4\} \cup \{R_3, R_4\}.$$

Any two elements of this set with distinct subscripts induce incompatible partitions of the set of marked points of D_n. As previously, it follows that every pair of subsurfaces in the set

$$\mathcal{X} = \{A_1, L_2 \cup A_2, L_3 \cup R_3, L_4 \cup A_4 \cup R_4\}$$

is mutually PB_n-incompatible. From this it follows that $\mathcal{X}$ is PB_n-overlapping. Applying Theorem 1.1 to all possible subsets of $\mathcal{X}$ (and choosing $\mathcal{F}$ appropriately) we obtain the desired conclusion (because each L_i and R_i that appears in X is either a disk with three or more marked points or an annulus with two or more marked points, they all can be realized as the support of a partial pseudo-Anosov element of PB_n). □

The situation for the Torelli groups is similar to that for the pure braid groups. Let X be a subsurface of S_g that is either a separating annulus or a subsurface with connected boundary.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

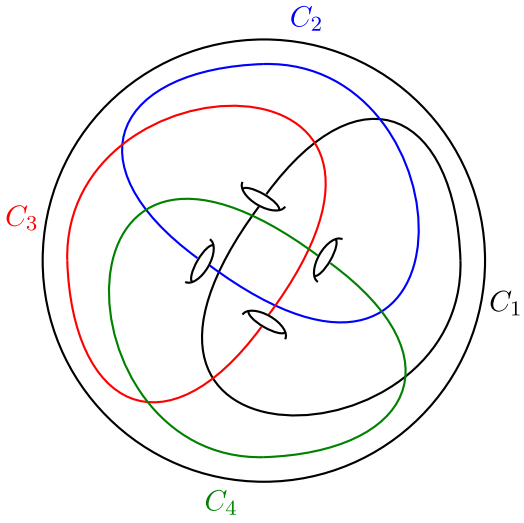

Figure 5. (Colour online) The curves in the proof of Theorem 2.6. Each C_i is symmetric about the plane of the page, and only half of each curve is shown.

In either case X induces a direct sum decomposition of $H_1(S_g;\mathbb{Q})$ into two symplectic subspaces. In the annulus case, the two subspaces are the elements of $H_1(S_g;\mathbb{Q})$ with representatives on one side or the other of the annulus. In the other case, the two subspaces are the elements of $H_1(S_g;\mathbb{Q})$ either represented in X or its complement. Similar to the case of the pure braid group, we say that two direct sum decompositions of $H_1(S_g;\mathbb{Q})$ are *compatible* if each summand of one decomposition either contains or is contained in a summand of the other decomposition.

In the following, the *genus* of a separating curve C in S_g is the minimum genus of a complementary component.

Theorem 2.6. *Suppose $g \geq 4$. For each of the 15 non-empty subsets A of $\{(1,0),(1,1),(0,2),(1,2)\}$ there is a normal subgroup of $\mathcal{I}(S_g)$ isomorphic to*

$$\underset{(a,b)\in A}{*}\Big(\underset{\infty}{*}\left((F_\infty)^b \times \mathbb{Z}^a\right)\Big).$$

Proof. Let C_1, C_2, C_3, and C_4 be four separating curves in S_g, each with genus at least two, and so that the corresponding direct sum decompositions of $H_1(S_g;\mathbb{Q})$ are pairwise incompatible. One such configuration for a surface of genus four is indicated in Figure 5; for higher genus we may add handles to any complementary region. We then define A_i, L_i, and R_i as in the proof of Theorem 2.3. (For the surface of genus four, any four curves of genus two giving distinct direct sum decompositions of $H_1(S_g;\mathbb{Q})$ will suffice, and any such configuration can be extended to higher genus by adding extra handles to a single complementary region).

Consider the following set of subsurfaces of S_g:

$$\{A_1, A_2, A_4\} \cup \{L_2, L_3, L_4\} \cup \{R_3, R_4\}.$$

The decomposition of $H_1(S_g;\mathbb{Q})$ induced by a separating curve in S_g is the same as the decomposition induced by either of its complementary regions. Thus, any two elements of the above set with distinct subscripts induce incompatible direct sum decompositions of $H_1(S_g;\mathbb{Q})$. Because $\mathcal{I}(S_g)$

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

acts trivially on $H_1(S_g;\mathbb{Q})$, it follows that every pair of subsurfaces in the set

$$\mathcal{X} = \{A_1, L_2 \cup A_2, L_3 \cup R_3, L_4 \cup A_4 \cup R_4\}$$

is mutually $\mathcal{I}(S_g)$-incompatible. It follows that $\mathcal{X}$ is $\mathcal{I}(S_g)$-overlapping. Applying Theorem 1.1 to all possible subsets of $\mathcal{X}$, and choosing $\mathcal{F}$ appropriately, we obtain the desired conclusion (we may apply, for example, the Thurston construction of pseudo-Anosov mapping classes to produce partial pseudo-Anosov elements of $\mathcal{I}(S_g)$ with support in any L_i or R_i). □

3. Projection complexes

In § 3.1, we recall the definition of a projection complex and explain how any tree may be viewed as a projection complex. Then, in § 3.2, we recall the theory of WPD elements and explain the construction of a projection complex from a collection of WPD elements. Finally, in § 3.3 we explain how to build a projection complex from the curve complexes of orbit overlapping subsurfaces of a given surface; our new contribution, Proposition 3.2, is a version for disconnected subsurfaces.

3.1 The definition

Let $\mathbb{Y}$ be a set and let $\theta \geq 0$ be a constant. Assume that for each $y \in \mathbb{Y}$ there is a function

$$d_y \colon (\mathbb{Y} \setminus \{y\}) \times (\mathbb{Y} \setminus \{y\}) \to \mathbb{R}_{\geq 0}$$

with the following properties.

- *Symmetry:* $d_y(x,z) = d_y(z,x)$ for all $x, y, z \in \mathbb{Y}$.
- *Triangle inequality:* $d_y(x,z) + d_y(z,w) \geq d_y(x,w)$ for all $x, y, z, w \in \mathbb{Y}$.
- *Inequality on triples:* $\min\{d_y(x,z), d_z(x,y)\} \leq \theta$ for all $x, y, z \in \mathbb{Y}$.
- *Finiteness:* $\#\{y \in \mathbb{Y} \mid d_y(x,z) > \theta\}$ is finite for all $x, z \in \mathbb{Y}$.

These conditions are known as the *projection complex axioms*. When we say that a set $\mathbb{Y}$ and a collection of such functions $\{d_y\}_{y\in\mathbb{Y}}$ satisfy the projection complex axioms the constant θ is implicit.

For a given $K \geq 0$, we define a graph $\mathcal{P}_K(\mathbb{Y})$ with vertices corresponding to the elements in $\mathbb{Y}$. To define the edges, we require the notion of modified distance functions.

Given the functions $\{d_y\}$, Bestvina–Bromberg–Fujiwara [BBF15] constructed another collection of functions $\{d'_y\}_{y\in\mathbb{Y}}$, where each d'_y shares the same domain and target as d_y. Because the definition of the d'_y is technical and because we do not use the definition in this paper, we do not state it here. Bestvina–Bromberg–Fujiwara [BBF15, Theorem 3.3B] showed that the modified functions are coarsely equivalent to the original functions: for $x \neq y \neq z \in \mathbb{Y}$, $d'_y(x,z) \leq d_y(x,z) \leq d'_y(x,z) + 2\theta$.

Fix $K \geq 0$. Then two vertices x, z of $\mathcal{P}_K(\mathbb{Y})$ are connected by an edge if $d'_y(x,z) \leq K$ for all $y \in \mathbb{Y} - \{x,z\}$. Let d denote the resulting path metric on $\mathcal{P}_K(\mathbb{Y})$.

Bestvina–Bromberg–Fujiwara showed that for K large enough relative to θ, there are constants C_e, C_p, and C_g, so that the following properties hold (see [BBF15, Proposition 3.14, and Lemma 3.18]).

(i) *Bounded edge image.* If $x \neq y \neq z$ are vertices of $\mathcal{P}_K(\mathbb{Y})$ and $d(x,z) = 1$, then $d_y(x,z) \leq C_e$.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

(ii) *Bounded path image.* If a path in $\mathcal{P}_K(\mathbb{Y})$ connects vertices x to z without passing through the 2-neighborhood of the vertex y, then $d_y(x,z) \leq C_p$.
(iii) *Bounded geodesic image.* If a geodesic in $\mathcal{P}_K(\mathbb{Y})$ connects vertices x to z without passing through the vertex y, then $d_y(x,z) \leq C_g$.

(The bounded edge image property follows from the definition of the edges of $\mathcal{P}_K(\mathbb{Y})$, with $C_e = K + 2\theta$.) If K is large enough so that the graph $\mathcal{P}_K(\mathbb{Y})$ satisfies the bounded edge, path, and geodesic properties for some C_e, C_p, and C_g, then we say that $\mathcal{P}_K(\mathbb{Y})$ is a *projection complex*.

For a projection complex, we refer to C_e, C_p, and C_g as the *edge constant*, the *path constant*, and the *geodesic constant*, respectively.

We note that our terminology is not standard; in the papers by Bestvina–Bromberg–Fujiwara [BBF15] and Bestvina–Bromberg–Fujiwara–Sisto [BBFS19], every $\mathcal{P}_K(\mathbb{Y})$ is called a projection complex.

Group actions on projection complexes. We say that a group G acts on a projection complex $\mathcal{P}_K(\mathbb{Y})$ if G acts on the set $\mathbb{Y}$ in such a way that the associated distance functions d_y are G-invariant, that is, $d_{gy}(gx, gz) = d_y(x,z)$. We note that if the original distance functions d_y are G-invariant, then the modified distance functions are G-invariant as well, and so the action of G on $\mathbb{Y}$ in particular extends to an action of G on the graph $\mathcal{P}_K(\mathbb{Y})$ by simplicial automorphisms.

Trees are projection complexes. As an illustration of the definition of a projection complex, we now explain how an arbitrary simplicial tree can be viewed as a projection complex. Let T be a simplicial tree and $\mathbb{Y}$ the set of vertices in T. We set

$$d_y(x,z) = \begin{cases} 1 & \text{if } y \text{ is on the geodesic from } x \text{ to } z, \\ 0 & \text{otherwise.} \end{cases}$$

These functions satisfy the projection complex axioms for any $\theta \geq 0$. Note that the inequality on triples axiom is merely the fact that for any triple of points in a tree, at most one is on the geodesic between the other two. For this example, the modified distance functions are the same as the original distance functions and moreover for any $0 < K < 1$ we see that $\mathcal{P}_K(\mathbb{Y})$ is isomorphic to T. On the other hand, if $K \geq 1$, then $\mathcal{P}_K(\mathbb{Y})$ is a complete graph.

3.2 Projection complexes from WPD elements

In this section, we give the first important general construction of a projection complex. We begin by briefly describing two motivating examples. Suppose we have either:

(i) a collection of elements of a free group acting on the Cayley graph of the free group; or
(ii) a collection of elements of a surface group acting on the hyperbolic plane.

In either case, there is a projection complex where the vertices are the orbits of the axes of the given elements and the distance functions are given as follows: if x, y, and z are axes for the given elements, then $d_y(x,z)$ is the diameter of the union of the nearest-point projections of x and z to y. It is an exercise in either graph theory or hyperbolic geometry to show that these functions give rise to a projection complex.

In the rest of this section, we describe a general setup for constructing projection complexes, inspired by these two examples. The end result is Proposition 3.1. In order to state it we require

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

a number of definitions. In the remainder of this section, let X be a hyperbolic metric space and let G be a group acting on X.

Distance functions from nearest-point projection. For a subset Z of X and a point $p \in X$, the *projection of p to Z* is the (possibly empty) subset

$$\pi_Z(x) = \{z \in Z \mid d(x,z) \leq d(x,z') \text{ for all } z' \in Z\}.$$

This defines a function $\pi_Z \colon X \to \wp(Z)$.

Let $\mathbb{Y}$ be a collection of subsets of X and for each a in $\mathbb{Y}$, define $\pi_a \colon X \to \wp(a)$ as previously. If $\pi_a(b)$ is bounded and non-empty for all pairs $a, b \in \mathbb{Y}$, we may define $d_a \colon \mathbb{Y} \setminus \{a\} \times \mathbb{Y} \setminus \{a\} \to \mathbb{R}$ by

$$d_a(b,c) = \operatorname{diam}(\pi_a(b) \cup \pi_a(c)).$$

Then the family $\{d_a\}_{a \in \mathbb{Y}}$ satisfies the first two projection complex axioms. Our goal in what follows is to describe conditions under which the other two axioms are also satisfied.

One easy special case is where X is a simplicial tree and $\mathbb{Y}$ is any collection of disjoint geodesics. To verify the last two axioms, the key observation is that the geodesics y with non-zero $d_y(x,z)$ are precisely those sharing an edge with the shortest path between x and z.

Translation length and hyperbolic elements. Let $f \in G$. The *translation length* of f is the limit

$$\tau(f) = \lim_{n \to \infty} \frac{d(p, f^n p)}{n}.$$

for some $p \in X$; this limit is independent of the point p. We say f is *hyperbolic* if $\tau(f) > 0$.

Weak proper discontinuity. We say $f \in G$ is a *WPD element* if it is hyperbolic and for all $D \geq 0$ and $p \in X$ there exists $M \geq 0$ such that the set

$$\{g \in G \mid d\big(p, gp\big) \leq D \text{ and } d\big(f^M p, g\big(f^M p\big)\big) \leq D\}$$

is finite. The WPD elements for the $\operatorname{Mod}(S)$ action on $\mathcal{C}(S)$ are precisely the pseudo-Anosov mapping classes [BF02].

Elementary closure. Suppose $f \in G$ is a WPD element and let $\mathcal{O}$ be the orbit of some point X under $\langle f \rangle$. Define the *elementary closure*, $\operatorname{EC}(f)$, to be the subgroup of $g \in G$ such that there exists a $\Delta \geq 0$ with $g\mathcal{O}$ contained in the Δ-neighborhood of $\mathcal{O}$. The group $\operatorname{EC}(f)$ is well-defined independent of the orbit. Bestvina–Bromberg–Fujiwara proved [BBF13, Proposition 4.7] that for such an $f \in G$, there is a short exact sequence

$$1 \to \operatorname{EC}_0(f) \to \operatorname{EC}(f) \to \Gamma \to 1,$$

where Γ is isomorphic to either $\mathbb{Z}$ or $\mathbb{Z}_2 * \mathbb{Z}_2$ and $\operatorname{EC}_0(f)$ is finite. The set $\operatorname{EC}(f) \cdot \mathcal{O}$ is called a *quasi-axis bundle* for f. It follows from the definitions that the stabilizer of any quasi-axis bundle is $\operatorname{EC}(f)$.

As X is hyperbolic, any hyperbolic isometry acts with north–south dynamics on the boundary of X. It follows readily that for a WPD element $f \in G$, the group $\operatorname{EC}(f)$ is the stabilizer in G of the pair of fixed points in ∂X for f. In particular, if $f \in \operatorname{Mod}(S)$ is pseudo-Anosov and X

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

is the curve complex $\mathcal{C}(S)$, then $\mathrm{EC}(f)$ is the stabilizer in G of the pair of projective measured foliations on S fixed by f, as defined in the introduction.

Normal independence. We say that WPD elements f_1 and f_2 in G are *independent* if $\mathrm{EC}(f_1) \cap \mathrm{EC}(f_2)$ is finite. This is equivalent to requiring that the fixed points sets for f_1 and f_2 in ∂X are disjoint. We further say that f_1 and f_2 are *normally independent* if every conjugate of f_1 is independent of f_2. We say that a collection of WPD elements is *(normally) independent* if they are pairwise (normally) independent.

For Proposition 3.1, we use the easy observation that normally independent elements cannot have conjugate powers. In fact, one can prove that two WPD elements of a group acting on a hyperbolic space are normally independent if and only if they fail to have conjugate non-trivial powers.

The following result of Dahmani–Guirardel–Osin [DGO17] will be used to construct a projection complex in § 7 (for the action of the mapping class group on the curve complex).

Proposition 3.1 (Dahmani–Guirardel–Osin). *Let G be a group and X a hyperbolic metric space on which G acts. Let $\{f_1, \dots, f_m\}$ be a normally independent family of WPD elements and for each f_i let $\beta_i \subseteq X$ be a quasi-axis bundle. Then the set $\mathbb{Y} = G \cdot \{\beta_1, \dots, \beta_m\}$ together with the distance functions $\{d_\beta\}_{\beta \in \mathbb{Y}}$ satisfy the projection complex axioms.*

If $\mathcal{P}$ is a projection complex arising from Proposition 3.1, and if the translation distances of the f_i are all bounded below by L, then the collection $\{\langle g f_i g^{-1} \rangle \mid g \in G,\ 1 \leq i \leq m\}$ is an equivariant L-spinning family of subgroups. We use this fact in the proof of Theorem 1.7 in § 7.

Dahmani–Guirardel–Osin did not state Proposition 3.1 explicitly, but it follows from their work as we now explain. Proposition 3.1 has the same hypothesis as Theorem 6.8 in their paper, whose proof hinges on showing that the groups $\mathrm{EC}(f_i)$ are geometrically separated, as defined in that paper. Lemma 4.47 in their paper proves, in particular, that the G-translates of the quasi-axis bundles (i.e. the orbits of the geometrically separated subgroups) and the associated distance functions satisfy the projection complex axioms. In their paper, the quasi-axes bundles β_i are all orbits of a common basepoint, but standard coarse geometry for hyperbolic spaces can be used to extend Lemma 4.47 to arbitrarily chosen quasi-axis bundles, up to changing the constants of the projection complex.

Proposition 3.1 actually holds with the assumption that X is a hyperbolic space replaced by the requirement that the WPD elements have strongly contracting axes. The argument in the case of a single element appears in the first version of a paper by Bestvina–Bromberg–Fujiwara [BBF13], and we prove the general case in a forthcoming paper.

3.3 Projection complexes from subsurface projections

We recall the Bestvina–Bromberg–Fujiwara construction of the projection complex associated to the curve complexes of a collection of orbit-overlapping connected subsurfaces of a given surface. We then give our generalization to the case of disconnected subsurfaces.

Projection complexes from connected G-overlapping subsurfaces. Given a non-annular connected subsurface $X \subseteq S$, and a curve γ in the curve complex $\mathcal{C}(S)$, the *projection* of γ to X is the subset $\pi_X(\gamma) \subset \mathcal{C}(X)$ of curves that have a representative disjoint from some arc in $X \cap \gamma$. When $X \subseteq S$ is an annulus, the definitions of $\mathcal{C}(X)$ and $\pi_X \colon \mathcal{C}(S) \to \wp(\mathcal{C}(X))$ are more complicated; see [MM00] for details. In either case, when $\pi_X(\gamma)$ is non-empty, the diameter of $\pi_X(\gamma)$ is at

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

most two [MM00, Lemma 2.3]. Moreover, the set $\pi_X(\gamma)$ is non-empty whenever γ essentially intersects X.

If $\mathbb{Y}$ is a $\mathrm{Mod}(S)$-invariant family of isotopy classes of connected subsurfaces of S such that $\pi_Y(\partial X)$ is non-empty for all distinct $X, Y \in \mathbb{Y}$, then subsurface projection defines a function $d_Y\colon \mathbb{Y} - \{Y\} \times \mathbb{Y} - \{Y\} \to \mathbb{R}_{\geq 0}$ by

$$d_Y(X, Z) = \mathrm{diam}_{\mathcal{C}(Y)}(\pi_Y(\partial X) \cup \pi_Y(\partial Z)).$$

These functions clearly satisfy the first two projection complex axioms. That the inequality on triples axiom is satisfied for $\theta = 10$ is known as the Behrstock inequality [Beh06, Man10]. Bestvina–Bromberg–Fujiwara [BBF15, Lemma 5.3] show that the finiteness axiom holds for $\theta = 3$.

Projection complexes from arbitrary G-overlapping subsurfaces. Let G be a subgroup of $\mathrm{Mod}(S)$ and let $\mathcal{X}$ be an G-overlapping family of subsurfaces of S. Set $\mathbb{Y} = G \cdot \mathcal{X}$. Given $Y \in \mathbb{Y}$ we define a function $\underline{d}_Y\colon \mathbb{Y} - \{Y\} \times \mathbb{Y} - \{Y\} \to \mathbb{R}_{\geq 0}$ by

$$\underline{d}_Y(X, Z) = \sum_{Y' \in \pi_0(Y)} d_{Y'}(X, Z),$$

where $d_{Y'}$ is the subsurface projection distance defined above. As $\mathcal{X}$ is an G-overlapping family, each term in the summand is defined. The following proposition gives us the projection complex used for Theorem 1.1.

Proposition 3.2. *Suppose G is a subgroup of $\mathrm{Mod}(S)$ and $\mathcal{X}$ is a G-overlapping family of subsurfaces. Then the set $\mathbb{Y} = G \cdot \mathcal{X}$ together with the distance functions $\{\underline{d}_Y\}_{Y \in \mathbb{Y}}$ satisfy the projection complex axioms.*

Proof. The first two projection complex axioms hold for the functions $\underline{d}_Y$, because they hold for the summands.

Let C denote the maximum number of components of a subsurface in $\mathbb{Y}$. Suppose $\underline{d}_Y(X, Z) > 12C$. Thus, $d_{Y'}(X, Z) > 12$ for some $Y' \in \pi_0(Y)$ and, hence, $d_{Y'}(X', Z') > 10$ for each $X' \in \pi_0(X)$ and $Z' \in \pi_0(Z)$. Thus, by the Behrstock inequality, we have $d_{X'}(Y, Z) \leq 10$ for any component $X' \in \pi_0(X)$. Hence, $\underline{d}_X(Y, Z) \leq 10C$. This shows that the inequality on triples axiom holds for $\{\underline{d}_Y\}_{Y \in \mathbb{Y}}$ using $\theta = 12C$.

Finally, as $\#\{Y \mid d_Y(X, Z) > 3\} < \infty$ for any connected subsurfaces $X, Z \subseteq S$, we see the same is true for $\{Y \in \mathbb{Y} \mid \underline{d}_Y(X, Z) > 5C\}$ for any $X, Z \in \mathbb{Y}$. Indeed, fix $X, Z \in \mathbb{Y}$. If $\underline{d}_Y(X, Z) > 5C$, then $d_{Y'}(X, Z) > 5$ for some $Y' \in \pi_0(Y)$. Hence, $d_{Y'}(X', Z') > 3$ for any $X' \in \pi_0(X)$ and $Z' \in \pi_0(Z)$. Thus, there are only finitely many possibilities for the subsurface Y'. However, as the subsurfaces in $\mathbb{Y}$ are either equal or overlapping, such a subsurface Y' can be a component for only a single subsurface in $\mathbb{Y}$. Thus, the finiteness axiom holds and the proposition is proved. □

4. Windmills and waypoints in projection complexes

Here we develop our theory of windmills for group actions on projection complexes. In § 6 we use this to prove Theorem 1.6, which states that if we have a group G acting on a projection complex $\mathcal{P}$ and if we have an equivariant L-spinning family $\{R_v\}$ of subgroups of G, then when L is large, the subgroup of G generated by all R_v is isomorphic to the free product of some of the R_v.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

We begin with the basic definitions regarding windmills, which allow us to rephrase Theorem 1.6 as Proposition 4.1. Then we introduce the notions of pivot points and waypoints, which gives a further reduction of Proposition 4.1, given in Proposition 4.2.

4.1 Windmill data

Given an action of a group G on a projection complex $\mathcal{P}$ with an equivariant family of subgroups $\{R_v\}$ of G, we can inductively define a sequence of subgraphs W_i of $\mathcal{P}$, a sequence of subsets $\mathcal{O}_i$ of the set of vertices of $\mathcal{P}$, and a sequence of subgroups H_i of G as follows. (In fact, the constructions of this section apply to group actions on arbitrary graphs.)

Let v_0 be some base point for $\mathcal{P}$. To begin the inductive definitions at $i = 0$, we define:

- $H_0 = R_{v_0}$; and
- $W_0 = \mathcal{O}_0 = \{v_0\}$.

For $i \geq 1$, we denote by N_i the 1-neighborhood of W_{i-1}, we denote by L_i the vertices of $N_i \setminus W_{i-1}$, and we define:

- $H_i = \langle R_v \mid v \in N_i \rangle$;
- $W_i = H_i \cdot N_i$; and
- $\mathcal{O}_i =$ a set of orbit representatives for the action of H_{i-1} on L_i.

The letter W here stands for 'windmill' and we refer to the set

$$\{(H_i, W_i, \mathcal{O}_i)\}_{i=0}^{\infty}$$

as a set of *windmill data* for the equivariant family $\{R_v\}$.

We can see inductively that each W_i is connected. The keys to the inductive step are that H_i acts on W_i and that each generator g of H_i has the property that $\big(gN_i\big) \cap N_i \neq \emptyset$.

Given this windmill data, we may further define both a sequence of groups and a sequence of homomorphisms, for $i \geq 0$:

- $F_i = \Big(\mathop{\ast}_{v \,\in\, \mathcal{O}_i} R_v \Big) * F_{i-1}$; and
- $\rho_i \colon F_i \to H_i$, the natural map.

In the definition of F_0 we treat F_{-1} as the trivial subgroup, and so $F_0 = R_{v_0}$. Note that, by definition, each ρ_i is injective on each $R_v \subseteq F_i$.

We refer to the collection $\{(F_i, \rho_i)\}_{i=0}^{\infty}$ as the *free product data* associated to the windmill data $\{(H_i, W_i, \mathcal{O}_i)\}_{i=0}^{\infty}$.

We denote the direct limit of the F_i by F:

$$F = \varinjlim F_i.$$

There is a natural partition of F into sets $\{F^{(i)}\}_{i=-1}^{\infty}$ defined by

$$F^{(i)} = F_i \setminus F_{i-1},$$

where we treat F_{-2} as the empty set, and so $F^{(-1)} = \{\mathrm{id}\}$ and $F^{(0)} = R_{v_0} \setminus \{\mathrm{id}\}$. We refer to each $F^{(i)}$ as the ith *level* of F.

The subgroup H of G generated by the R_v is the direct limit of the H_i. With this setup in hand we may rephrase Theorem 1.6 in the following way.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Proposition 4.1. *Let $\mathcal{P}$ be a projection complex, G a group that acts on $\mathcal{P}$, and $\{R_v\}$ an equivariant family of subgroups. Choose windmill data $\{(H_i, W_i, \mathcal{O}_i)\}_{i=0}^{\infty}$. Then the conclusion of Theorem 1.6 holds if each ρ_i is injective.*

Proof. If v and w are two vertices of $\mathcal{P}$ in the same H-orbit, then by the equivariance condition, we have $R_v \cong R_w$. Thus, we may check that the conclusion of Theorem 1.6 holds for any particular choice $\mathcal{O}$ of orbit representatives. We check that it holds for $\mathcal{O} = \bigcup \mathcal{O}_i$.

First we verify that $\bigcup \mathcal{O}_i$ is indeed a set of orbit representatives. The equivariance condition means two vertices v, w in the same orbit have conjugate corresponding subgroups R_v and R_w. On the other hand, if $v, w \in \bigcup \mathcal{O}_i$, then R_v and R_w are free factors in the domain F_i of ρ_i, and distinct free factors of a free product have trivially intersecting conjugates. Thus, injectivity of the ρ_i implies that distinct vertices of $\bigcup \mathcal{O}_i$ represent distinct H-orbit representatives. We must also show that every H-orbit is represented in $\bigcup \mathcal{O}_i$. Given O an H-orbit of vertices in $\mathcal{P}$, there exists a minimal i such that $O \cap N_i$ is not empty. By minimality of i, if w is in $O \cap N_i$, then it is not contained in $W_{i-1} = H_{i-1} \cdot N_{i-1}$. In particular, $w \in L_i$ and, thus, has an orbit representative in $\mathcal{O}_i$.

The ρ_i induce a natural surjective map $\rho\colon F \to H$. The statement that $\langle R_v \rangle$ is isomorphic to the free product

$$\ast_{v \in \bigcup \mathcal{O}_i} R_v$$

is equivalent to the statement that ρ is injective, which is in turn equivalent to the statement that each ρ_i is injective. This is true by hypothesis, and so the proof is complete. □

4.2 Pivot points and waypoints

We continue with the notation of the previous section. To each element of F at a given level $i \geq 0$, we associate a subset of the vertices of $\mathcal{P}$, as follows. Each $h \in F^{(i)}$ has a *syllable decomposition* $h_1 \cdots h_n$ with respect to the free product decomposition used to define F_i. Specifically, each syllable h_j is either a non-trivial element of F_{i-1} or a non-trivial element of R_{v_j} with $v_j \in \mathcal{O}_i$, and also the following property is satisfied: no two consecutive syllables are of the first type and consecutive syllables h_j and h_{j+1} of the second type have distinct corresponding fixed vertices v_j and v_{j+1}. We refer to n as the *syllable length* of h.

Let $i \geq 1$ and fix some $h \in F^{(i)}$ with syllable decomposition $h = h_1 \cdots h_n$. For $j \in \{1, \ldots, n\}$ with $h_j \notin F_{i-1}$ and with corresponding fixed vertex v_j we define a vertex w_j of $\mathcal{P}$ as follows:

$$w_j = h_1 \ldots h_{j-1} v_j.$$

Note that v_j and w_j are not defined for $h_j \in F_{i-1}$. Let $\mathrm{Piv}(h)$ be the ordered list of points w_j, and call these the *pivot points* for h. For $h \in F_0$ we define $\mathrm{Piv}(h)$ to be empty.

Each group F_i acts on $\mathcal{P}$ via ρ_i. We say that a vertex v of $\mathcal{P}$ is a *waypoint* for $h \in F_i$ if it is not equal to v_0 or hv_0 and every geodesic from v_0 to hv_0 passes through v. If at least one pivot point for h is a waypoint, it follows that $\rho_i(h)$ is non-trivial. Recall that Proposition 4.1 reduces Theorem 1.6 to proving injectivity of ρ_i for all i. Hence, we have further reduced our task to proving that every non-trivial $h \in F^{(i)}$ has a waypoint, for all $i \geq 1$. We record this fact in the following proposition.

Proposition 4.2. *Let $\mathcal{P}$ be a projection complex, G a group that acts on $\mathcal{P}$, and $\{R_v\}$ an equivariant family of subgroups. Choose windmill data $\{(H_i, W_i, \mathcal{O}_i)\}_{i=0}^{\infty}$ and let $\{(F_i, \rho_i)\}_{i=0}^{\infty}$*

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

be the associated free product data. If for each $i \geq 1$ and each $h \in F^{(i)}$ at least one element of $\mathrm{Piv}(h)$ *is a waypoint for h, then the conclusion of Theorem 1.6 holds.*

5. Windmills in trees

In this section, we show how the windmill machinery developed in § 4 is used in a special case, namely, the case of a group acting on a tree. Let G be a group acting on a tree $\mathcal{T}$, with $\{R_v\}$ an equivariant spinning family of subgroups. The spinning condition takes a particularly simple form in the case of a tree, as follows.

- *Spinning condition for trees:* For any vertex v of $\mathcal{T}$, any incident edge e of $\mathcal{T}$, and any non-trivial $h \in R_v$ we have
$$h(e) \neq e.$$

We have the following special case of Theorem 1.6.

Theorem 5.1. *Let $\mathcal{T}$ be a tree and let G be a group acting on $\mathcal{T}$. If $\{R_v\}$ is an equivariant spinning family of subgroups of G with windmill data $\{(H_i, W_i, \mathcal{O}_i)\}_{i=0}^{\infty}$, then the subgroup H of G generated by the R_v is isomorphic to the free product*
$$\ast_{v \in \bigcup \mathcal{O}_i} R_v.$$

Theorem 5.1 is a recasting of a standard fact from the Bass–Serre theory of group actions on trees. As the R_v are vertex stabilizer subgroups, it follows from Bass–Serre theory that the quotient $\mathcal{T}/H$ is a tree and, hence, that H is isomorphic to a free product
$$\ast R_{v_j},$$
where $\{v_j\}$ is a set of orbit representatives for the action of H on the set of vertices of $\mathcal{T}$ (see [Ser03, § 5.1]). Because $\mathcal{T}$ is a tree, we can find such a set of representatives by choosing a lift of $\mathcal{T}/H$ to $\mathcal{T}$. The windmill data, more specifically $\mathcal{O}_i$, exactly describe such a set of representatives, as shown in the proof of Proposition 4.1.

The proof that we give of Theorem 5.1 can be thought of as a rephrasing of the standard argument from Bass–Serre theory. This rephrasing models our proof of the more general Theorem 1.6. Before giving the proof, we require the following technical lemma, which is also required for the proof of Theorem 1.6.

A consequence of distinctness of waypoints. In the proof of Theorems 5.1 and 1.6 we show inductively that the set of pivot points $\mathrm{Piv}(h)$ for an element $h \in F^{(i)}$ are distinct waypoints. The inductive argument uses the following technical lemma.

Lemma 5.2. *Let $\mathcal{P}$ be a projection complex, G a group that acts on $\mathcal{P}$, and $\{R_v\}$ an equivariant family of subgroups. Choose windmill data and let $\{(F_i, \rho_i)\}_{i=0}^{\infty}$ be the associated free product data. Fix $i \geq 0$, and suppose the pivot points* $\mathrm{Piv}(h)$ *are distinct waypoints for each $h \in F_i$. If $h \in F_i$ and $\langle h \rangle$ has a bounded orbit, then h lies in R_w for some $w \in W_i$.*

Proof. Consider $h \in F_i$ at some level $j \leq i$. If $\langle h \rangle$ has a bounded orbit, then the number of waypoints in $\mathrm{Piv}(h^n)$ is bounded, independently of n, and therefore the length of the syllable decomposition of $h^n \in F^{(j)}$ is also bounded, independently of n. This means that h is conjugate

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

to an element with syllable length at most one. As $h \in F^{(j)}$, that means h lies in gR_vg^{-1} with $v \in \bigcup_{j=0}^{i} \mathcal{O}_j$ where $j \leq i$ and $g \in F_i$. Therefore h lies in R_w for $w = gv$. □

Proof of Theorem 5.1. As in the definition of the windmill data, W_0 consists of a basepoint, which we denote by v_0. In addition, for $i \geq 1$, we denote the 1-neighborhood of W_{i-1} in $\mathcal{T}$ by N_i and $L_i = N_i \setminus W_{i-1}$. Let $\{(F_i, \rho_i)\}_{i=0}^{\infty}$ be the free product data associated to the given windmill data, let F denote the direct limit of the F_i, let H denote the direct limit of the H_i, and let $F^{(i)}$ denote the ith level of F.

As per Proposition 4.2, it is enough to show for each $i \geq 1$ and each $h \in F^{(i)}$ that some element of $\mathrm{Piv}(h)$ is a waypoint. To prove this for all $i \geq 1$, we prove the following stronger statement by induction on $i \geq 0$: for each element $h \in F_i$, the pivot points of $\mathrm{Piv}(h)$ are waypoints and they appear in order along the geodesic from v_0 to hv_0 (in particular, they are distinct, which will allow us to apply Lemma 5.2). To mirror the proof of Theorem 1.6 in § 6, we break up the inductive hypothesis into two statements.

(A) If $h \in F_i$ and $w \in \mathrm{Piv}(h)$, then w is a waypoint for h.
(B) If $h \in F_i$, the elements of $\mathrm{Piv}(h)$ appear in order along the geodesic from v_0 to hv_0.

The base cases of statements (A) and (B), when $i = 0$, are vacuous. Thus, henceforth assume that $i \geq 1$, and that our inductive hypothesis holds for all h with level less than i. We now give the inductive steps for statements (A) and (B) in turn.

Proof of statement (A). It suffices to consider $h \in F^{(i)}$. We consider a secondary induction on the syllable length n of h. Thus, suppose that h has syllable decomposition $h_1 \cdots h_n$ where $n \geq 0$. For the secondary induction, our base case is $n = 0$. As $i \geq 1$ and the identity lies in $F^{(-1)}$, there is no element of $F^{(i)}$ with syllable length zero, and so the base case for the secondary induction is vacuous.

Assume now that $n \geq 1$. Let $w_k \in \mathrm{Piv}(h)$. The pivot point w_k corresponds to some syllable h_k that lies at level i. Let $h_\sigma = h_1 \cdots h_{k-1}$ and $h_\tau = h_{k+1} \cdots h_n$ (σ and τ stand for 'starting' and 'terminal'). Let v_k denote the vertex of $\mathcal{O}_i$ with $h_k \in R_{v_k}$. With this notation in hand, we write

$$h = h_\sigma h_k h_\tau \quad \text{and so} \quad w_k = h_\sigma v_k.$$

Applying h_σ^{-1} to all vertices, statement (A) is equivalent to the statement that v_k lies strictly between $h_\sigma^{-1} v_0$ and $h_k h_\tau v_0$.

Consider the diagram in Figure 6, which shows the case when $h_\sigma, h_\tau \in F^{(i)}$. There we have drawn the geodesics from v_0 to the vertices $h_\sigma^{-1} v_0$, v_k, and $h_\tau v_0$, respectively. These paths are solid, dashed, and dotted, respectively. The paths might overlap near v_0, as suggested by the figure. Let P be the geodesic from $h_\sigma^{-1} v_0$ to v_k and let Q be the geodesic from v_k to $h_\tau v_0$. As h_k fixes v_k, we have that $h_k Q$ is the geodesic path from v_k to $h_k h_\tau v_0$ and moreover that the concatenation of P and $h_k Q$ is a path from $h_\sigma^{-1} v_0$ to $h_k h_\tau v_0$ that contains v_k.

To prove statement (A), then, it is enough to prove that both P and Q are non-trivial and that the concatenation of P and $h_k Q$ is a geodesic. To do this, it is enough to show that both P and Q contain the (unique) edge e connecting v_k to W_{i-1}; equivalently that P and Q each contain a vertex of N_i other than v_k. It then follows from the spinning hypothesis that $h_k(e) \neq e$, and hence that the concatenation of P and $h_k Q$ is geodesic near v_k, hence is a geodesic. We treat the case of Q, with the case of P being the essentially the same.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

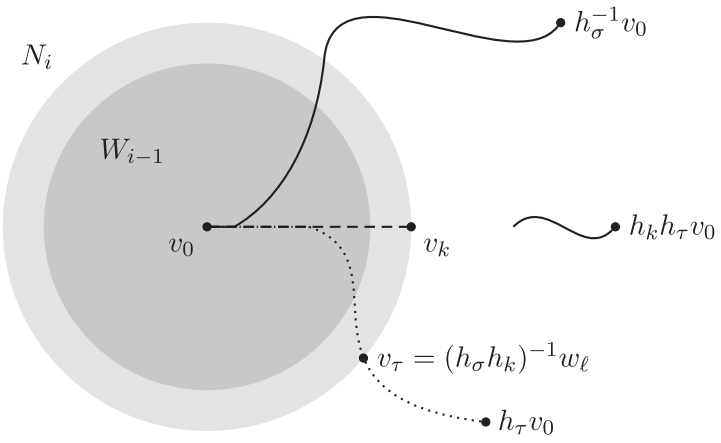

Figure 6. The vertices and paths used in the proof of Theorem 5.1.

There are two cases, according to whether h_τ lies in F_{i-1} or in $F^{(i)}$. In the first case, we have that $h_\tau v_0$ lies in W_{i-1}. As $h_\tau v_0$ is a vertex of Q, this completes the proof in the first case.

The second case is where h_τ lies in $F^{(i)}$. Let v_τ be the first pivot point for h_τ. We claim that v_τ lies on Q (as shown in Figure 6). Indeed, by the induction on syllable length applied to h_τ, we have that v_τ lies on the geodesic from v_0 to $h_\tau v_0$. Further, because v_τ is a leaf of the tree N_i, and because v_0 lies in N_i, any geodesic from a vertex of N_i to $h_\tau v_0$ contains v_τ, in particular the path Q contains v_τ (recall that Q contains $v_k \in N_i$). Thus, to show that Q contains a vertex of N_i other than v_k, it suffices to show the following claim.

Claim. The vertex v_τ is distinct from v_k.

From the definitions, v_τ is the fixed point of the first syllable of h_τ that is not contained in F_{i-1}. If this syllable is h_{k+1}, then by the definition of syllable decomposition,we have $v_k \neq v_{k+1} = v_\tau$.

The other case of the claim is where $h_{k+1} \in F_{i-1}$. In this case we must show that $v_k \neq h_{k+1}v_{k+2} = v_\tau$. Assume for the sake of contradiction that $v_k = h_{k+1}v_{k+2}$. To obtain the contradiction, we show that h_{k+1} fixes v_k, that is also lies in some R_w, and further that $w \neq v_k$. Together, these three items violate the spinning hypothesis.

As v_k and v_{k+2} are vertices of $\mathcal{O}_i$, which is a set of orbit representatives for the action of H_{i-1} on L_i, it follows from the assumption $v_k = h_{k+1}v_{k+2}$ that $v_k = v_{k+2}$. Thus, v_k is fixed by h_{k+1}, which is the first item. By the inductive hypothesis that all pivot points are waypoints for elements in F_{i-1} and because h_{k+1} fixes v_k, it follows from Lemma 5.2 that h_{k+1} lies in some R_w with $w \in W_{i-1}$, which is the second item. As v_k lies in L_i, and because the latter is disjoint from W_{i-1}, we obtain the third item. This completes the proof of the claim, hence statement (A).

Proof of statement (B). Take $h \in F^{(i)}$ as in the proof of statement (A). It suffices to show that if w_k is an element of $\mathrm{Piv}(h)$ that is not last in the order, and if w_ℓ is the immediately following element of $\mathrm{Piv}(h)$, then w_ℓ lies strictly between w_k and hv_0 in $\mathcal{T}$. As previously, we fix the decomposition $h_\sigma h_k h_\tau$ of h and we take v_k to be the vertex of $\mathcal{T}$ associated to h_k. We have that $w_\ell = h_\sigma h_k v_\tau$ where v_τ is the first pivot point for h_τ. We may assume that h_τ lies in $F^{(i)}$, for otherwise w_k would be the final pivot point. Similarly to our restatement of statement (A), we may apply $(h_\sigma h_k)^{-1}$ to all vertices in statement (B) in order to obtain the equivalent statement that v_τ lies strictly between v_k and $h_\tau v_0$. This is further equivalent to the statement that v_τ lies

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

in the interior of Q, which is the geodesic from v_k and $h_\tau v_0$. We observed this fact previously, and so the proof is complete. □

6. Proof of main technical theorem

In this section we prove our main technical theorem, Theorem 1.6. This theorem states that if a group G acts on a projection complex $\mathcal{P}$ with an equivariant L-spinning family of subgroups $\{R_v\}_{v\in\mathcal{P}}$ and L is sufficiently large, then the normal subgroup generated by $\{R_v\}_{v\in\mathcal{P}}$ is a free product of some of the R_v.

Our proof comes after a lemma that speaks to the difference between Theorem 1.6 and the special case considered in Theorem 5.1 about trees. The proofs of both theorems use windmills in virtually the same way, but the general case is complicated by the fact that windmills in a general projection complex need not be convex. Rather, we have Lemma 6.1, which is critical to our proof.

A convexity-like property of windmills in projection complexes. In the statement of the following lemma, C_p is the path constant associated to $\mathcal{P}$, as in § 3.1.

Lemma 6.1. *Let $\mathcal{P}$ be a projection complex, G a group that acts on $\mathcal{P}$, and $\{R_v\}$ an equivariant family of subgroups with windmill data $\{(H_i, W_i, \mathcal{O}_i)\}_{i=0}^{\infty}$. If x and y are vertices of W_{i-2} and $v \notin W_i$, then $d_v(x, y) \le C_p$.*

Proof. For $i \ge 2$, the distance between any vertex of $\mathcal{P} \setminus W_i$ and any vertex of W_{i-2} in $\mathcal{P}$ is at least three. This follows from the definitions, because each W_i contains the 1-neighborhood of W_{i-1}. As W_{i-2} is connected, the lemma now follows from the bounded path image property. □

Proof of Theorem 1.6. As in the statement, G is a group acting on a projection complex $\mathcal{P}$. Let θ be the constant from the definition of $\mathcal{P}$, and let C_e, C_g, and C_p be the associated constants from § 3.1, all of which ultimately depend on θ. We prove the theorem for $L(\mathcal{P}) = 3(11C_e + 6C_g + 5C_p + \theta) + 1$. In other words, we assume that G has an equivariant L-spinning family of subgroups $\{R_v\}$ with $L \ge L(\mathcal{P})$, and show that the normal closure of the union of the R_v can be decomposed into the free product described in the statement. As in § 4, let $\{H_i, W_i, \mathcal{O}_i\}$ be windmill data associated to the action of G on $\mathcal{P}$ and let $\{F_i, \rho_i\}$ be the associated free product data. As in the definition of the windmill data, W_0 consists of a basepoint, which we denote v_0. In addition for $i \ge 1$, we denote the 1-neighborhood of W_{i-1} in $\mathcal{P}$ by N_i and $L_i = N_i \setminus W_{i-1}$.

Let $m = 11C_e + 6C_g + 5C_p$; in particular, $L(\mathcal{P}) = 3(m + \theta) + 1$. To prove the theorem we show by induction on i that the following statements hold for all $i \ge 0$.

(A) If $h \in F_i$ and $w \in \mathrm{Piv}(h)$, then $d_w(v_0, hv_0) > m + \theta$.
(B) If $h \in F_i$, the elements of $\mathrm{Piv}(h)$ appear in order along any geodesic from v_0 to hv_0.
(C) For $x \in N_{i+1}$ and $v \notin W_i$ with $v \ne x$, we have $d_v(v_0, x) \le m$.

Respectively, these statements say that pivot points are waypoints, waypoints are distinct, and vertices outside a given W_i see bounded projection between points inside W_i (even N_{i+1}), which strengthens the convexity-like property of Lemma 6.1. The precise relevance of each statement to the proof is as follows. By the bounded geodesic image property, statement (A) implies that the pivot points of h are waypoints, because $m + \theta \ge C_g$. By Proposition 4.2, this will complete the proof of the theorem, once verified for all i. Statement (B) further implies that the waypoints

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

are distinct, which allow us to apply Lemma 5.2 in the same way as in the proof of Theorem 5.1. Statement (C) enables induction, as in effect it upgrades to W_i the W_{i-2} in Lemma 6.1.

For the base case of our induction on i, we consider the level $i = 0$. Statements (A) and (B) are vacuous. The hypotheses of statement (C) imply that $v \notin \{x, v_0\}$. As $d(v_0, x) = 1$, we have $d_v(v_0, x) \leq C_e \leq m$, where C_e is the edge constant from § 3.1.

Now consider the level $i \geq 1$ and assume the three statements hold for smaller levels. We prove each statement in turn. First, we prove statement (A) by our assumption on lower levels, and inducting on the syllable length of h. Then we upgrade statement (A) to statement (B). These are key to proving the new level of statement (C), which is, in turn, used to prove the next level of statements (A) and (B).

Proof of statement (A). It suffices to consider $h \in F^{(i)}$. We begin a second induction on the syllable length of h, using the decomposition $h = h_1 \cdots h_n$ described in § 4.2. For the base case of the induction on syllable length, suppose $n = 0$. As $i \geq 1$ and the identity lies in $F^{(-1)}$, there is no element of $F^{(i)}$ with syllable length zero, and so the base case for the secondary induction is vacuous.

Now suppose that $n \geq 1$ and that statement (A) holds for elements of $F^{(i)}$ with syllable length less than n. Let $w_k \in \mathrm{Piv}(h)$ be a pivot point for h, so $w_k = h_1 \cdots h_{k-1} v_k$ where $v_k \in \mathcal{O}_i$ is fixed by h_k. As in § 5, let $h = h_\sigma h_k h_\tau$, where h_σ and h_τ are possibly trivial subwords.

Step 1: the quantity $d_{w_k}(v_0, hv_0)$ is defined, as are $d_{v_k}(v_0, h_\sigma^{-1} v_0)$ and $d_{v_k}(v_0, h_\tau v_0)$. We use the latter two quantities in Step 2. The quantity $d_{w_k}(v_0, hv_0)$ is defined if and only if $w_k \notin \{v_0, hv_0\}$. We start by showing that $w_k \neq hv_0$. There are two cases: either $h_\tau \in F_{i-1}$ or $h_\tau \in F^{(i)}$. In the first case, we have $h_\tau v_0 \in F_{i-1} \cdot N_{i-1} = W_{i-1}$. As $v_k \in L_i$, which is disjoint from W_{i-1}, we have $v_k \neq h_\tau v_0$ and therefore $w_k = h_\sigma v_k = h_\sigma h_k v_k \neq hv_0$ as desired.

The second case is where h_τ lies in $F^{(i)}$. Let v_τ be the first pivot point for h_τ. By induction on n, statement (A) gives $d_{v_\tau}(v_0, h_\tau v_0) > m + \theta$. We observe that both v_τ and v_k are elements of L_i. We claim that $v_\tau \neq v_k$. This is exactly the claim from the proof of Theorem 5.1, whose proof uses the inductive hypothesis that statements (A) and (B) are true for the lower levels to invoke Lemma 5.2. In this context the same proof works after replacing 'spinning' with 'L-spinning'. Hence, we may apply statement (C) inductively (with respect to i) to conclude $d_{v_\tau}(v_0, v_k) \leq m$. As we have shown that $d_{v_\tau}(v_0, h_\tau v_0) > m + \theta$ and that $d_{v_\tau}(v_0, v_k) \leq m$, it follows that $v_k \neq h_\tau v_0$ and therefore, as in the first case, we have $w_k \neq hv_0$.

To confirm that $d_{w_k}(v_0, hv_0)$ is defined, it remains to show that $w_k \neq v_0$. To see this, we first observe that $h^{-1} w_k$ is a pivot point for h^{-1}. Hence, the previous argument shows that $h^{-1} w_k \neq h^{-1} v_0$ and, thus, $w_k \neq v_0$ as well.

To address the quantities $d_{v_k}(v_0, h_\sigma^{-1} v_0)$ and $d_{v_k}(v_0, h_\tau v_0)$, we observe that v_k is a pivot point for $h_k h_\tau$, which lies in $F^{(i)}$ and has syllable length at most n. Thus, applying the same argument to $h_k h_\tau$ instead of h, we deduce that $v_k \notin \{v_0, h_k h_\tau v_0\}$. In particular, $v_k \neq v_0$. Further, because v_k is fixed by h_k we also have $v_k \neq h_\tau v_0$. Similarly we have $v_k \notin \{v_0, h_\sigma^{-1} v_0\}$ as well.

Step 2: we have $d_{w_k}(v_0, hv_0) > m + \theta$. Using the invariance of the distance functions under the group action, the triangle inequality, and the L-spinning hypothesis, we have

$$\begin{aligned} d_{w_k}(v_0, hv_0) &= d_{v_k}(h_\sigma^{-1} v_0, h_k h_\tau v_0) \\ &\geq d_{v_k}(v_0, h_k v_0) - d_{v_k}(v_0, h_\sigma^{-1} v_0) - d_{v_k}(h_k v_0, h_k h_\tau v_0) \\ &\geq L - (d_{v_k}(v_0, h_\sigma^{-1} v_0) + d_{v_k}(v_0, h_\tau v_0)). \end{aligned}$$

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

We show that $d_{v_k}(v_0, h_\tau v_0) \leq m + \theta$. As h_σ^{-1} has functionally identical features (replacing h with h^{-1}), the same argument also shows that $d_{v_k}(v_0, h_\sigma^{-1} v_0) \leq m + \theta$. Incorporating this into above inequality, we conclude that

$$d_{w_k}(v_0, hv_0) \geq d_{v_k}(v_0, h_k v_0) - 2(m + \theta) \geq L - 2(m + \theta) > m + \theta, \tag{1}$$

which completes the proof of statement (A).

Towards bounding $d_{v_k}(v_0, h_\tau v_0)$, there are the same two cases as in Step 1: either $h_\tau \in F_{i-1}$ or $h_\tau \in F^{(i)}$. In the first case, we observed in Step 1 that $h_\tau v_0 \in W_{i-1} \subseteq N_i$, $v_k \notin W_{i-1}$ and that $v_k \neq h_\tau v_0$. Thus by induction on level i, statement (C) implies that $d_{v_k}(v_0, h_\tau v_0) \leq m \leq m + \theta$.

It remains to consider the second case. Let v_τ be the first pivot point in $\mathrm{Piv}(h_\tau)$. In this case, we have already observed in Step 1 that $d_{v_\tau}(v_0, h_\tau v_0) > m + \theta$ and $d_{v_\tau}(v_0, v_k) \leq m$. We thus have

$$d_{v_\tau}(v_k, h_\tau v_0) \geq d_{v_\tau}(v_0, h_\tau v_0) - d_{v_\tau}(v_0, v_k) > (m + \theta) - m > \theta.$$

Using the inequality on triples, we obtain $d_{v_k}(v_\tau, h_\tau v_0) \leq \theta$. As shown in Step 1, v_k and v_τ are distinct vertices of L_i, we apply statement (C) inductively (with respect to i) to obtain that $d_{v_k}(v_0, v_\tau) \leq m$. We thus have

$$d_{v_k}(v_0, h_\tau v_0) \leq d_{v_k}(v_\tau, h_\tau v_0) + d_{v_k}(v_0, v_\tau) \leq \theta + m,$$

completing the proof of statement (A).

Proof of statement (B). Recall that $\mathrm{Piv}(h) = \{w_j\}_{j \in I}$ is an ordered set with index set I a (possibly proper) subset of $\{1, \ldots, n\}$. Our goal is to show that these vertices appear in order on any geodesic from v_0 to hv_0.

Let $w_k, w_\ell \in \mathrm{Piv}(h)$ be consecutive, so $w_k = h_\sigma v_k$ and $w_\ell = h_\sigma h_k v_\tau$ where v_τ is the first pivot point of h_τ, again using the syllable decomposition $h = h_\sigma h_k h_\tau$. As previously, we have that $v_k \neq v_\tau$ and, thus, $w_k \neq w_\ell$. Using the triangle inequality, statement (A), and invariance of the distance functions, we have

$$d_{w_k}(v_0, w_\ell) \geq d_{w_k}(v_0, hv_0) - d_{w_k}(w_\ell, hv_0) \geq m + \theta - d_{v_k}(v_\tau, h_\tau v_0).$$

As in the proof of statement (A), using the inequality on triples, we may deduce that $d_{v_k}(v_\tau, h_\tau v_0) \leq \theta$ and, therefore, $d_{w_k}(v_0, w_\ell) \geq m > C_g$. By the bounded geodesic image property, any geodesic from v_0 to w_ℓ passes through w_k.

Proof of statement (C). Recall we are given $v \notin W_i$ and $x \in N_{i+1}$ such that $x \neq v$, and our goal is to bound $d_v(x, v_0)$ by $m = 11C_e + 6C_g + 5C_p$. There is an $x_0 \in W_i$ with $d(x, x_0) \leq 1$. If there is a geodesic from x_0 to v_0 that avoids v, in particular if W_i is convex, then by the bounded edge and geodesic image properties we have

$$d_v(v_0, x) \leq d_v(v_0, x_0) + d_v(x_0, x) \leq C_g + C_e \leq m.$$

However, it might be the case that every geodesic from x_0 to v_0 passes through v. In this case, we build a path from x_0 to v_0 consisting of: at most 10 edges avoiding v, at most six geodesics avoiding v, and at most five paths that each lie in some F_i-translate of W_{i-2}. As $v \notin W_i$ and W_i is F_i-invariant, the projection of any path in an F_i-translate of W_{i-2} to v is bounded by C_p by Lemma 6.1. Using this fact and the bounded edge and geodesic image properties, it follows that

$$d_v(v_0, x) \leq d_v(v_0, x_0) + d_v(x_0, x) \leq (10C_e + 6C_g + 5C_p) + C_e = m.$$

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

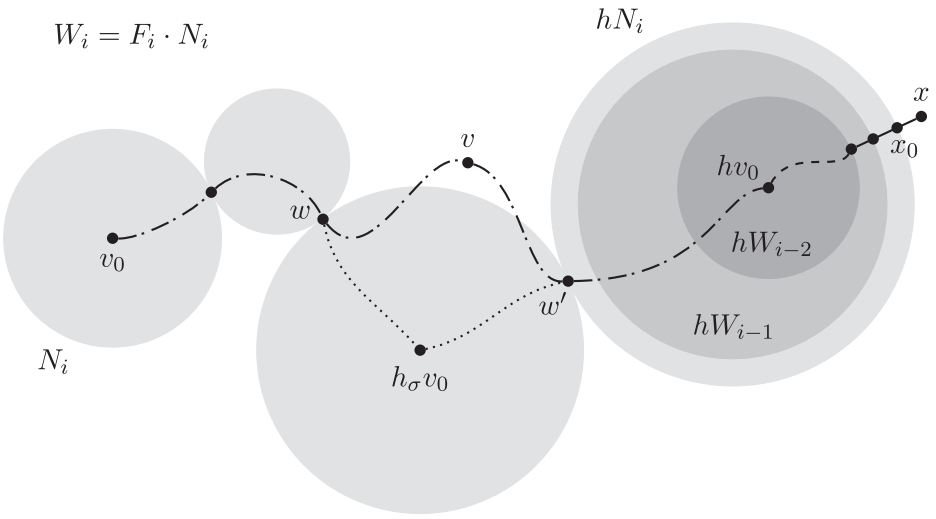

FIGURE 7. Steps 1 and 2 of statement (C) in the proof of Theorem 1.6.

Figure 7 illustrates the idea. In the rest of the proof we refer to edges avoiding v, geodesics avoiding v, and paths lying in an F_i-translate of W_{i-2} as paths of type e, g, and p, respectively. In addition, we assume throughout that $i > 1$; the case $i = 1$ follows the same idea but is simpler.

Step 1: connect x_0 to hv_0 for some $h \in F_i$. Using the definition of a windmill, there is an $h \in F_i$, $x' \in W_{i-2}$ and a path of length 2 from x_0 to hx' that does not include v (two paths of type e). As W_{i-2} is connected, there is a path in hW_{i-2} from hv_0 to hx' (path of type p).

Step 2: replace a segment of a geodesic from v_0 to hv_0 if it passes through v. If there is a geodesic from v_0 to hv_0 avoiding v we are done, since this is a path of type g. Otherwise, v lies on every geodesic from v_0 to hv_0. Fix such a geodesic. There are consecutive points $w, w' \in \mathrm{Piv}(h) \cup \{v_0, hv_0\}$ such that v lies between w and w' on this geodesic (extend the order on $\mathrm{Piv}(h)$ so v_0 is minimal and hv_0 is maximal). The (possibly trivial) subgeodesics from v_0 to w and from w' to hv_0 are of type g. By definition of the pivot points, w and w' are in $h_\sigma N_i$ for an initial subword h_σ of h. We replace the segment of the geodesic from w to w' by a concatenation of geodesics from w to $h_\sigma v_0$ and from $h_\sigma v_0$ to w'. If both of these geodesics avoid v (so they are of type g), then we are done. Suppose, in contrast, that one or both passes through v. We treat the case that the geodesic from w to $h_\sigma v_0$ passes through v; the other case is identical.

Step 3: connect w to $h_\sigma h' v_0$ for some $h' \in F_{i-1}$. As $w \in h_\sigma N_i$, this is identical to Step 1 with h, x_0, and hv_0 replaced by h_σ, w, and $h_\sigma h' v_0$. Like Step 1, this step contributes two paths of type e and at most one path of type p.

Step 4: replace a segment of a geodesic from $h_\sigma v_0$ to $h_\sigma h' v_0$ if it is passes through v. We proceed as in Step 2. The geodesic splits into two paths of type g and geodesic from u and u' that lie in $h_\sigma h'_\sigma N_{i-1}$ for some initial subword h'_σ (here u and u' are analogous to w and w' from Step 2). There is a path from u to u' consisting of two paths of type e and a path of type p. □

7. Proof of the pseudo-Anosov case

In this section, we prove Theorem 1.7, which states that if a group G acts on a hyperbolic space X and $\{f_1, \dots, f_m\} \subseteq G$ is a normally independent collection of NEC WPD elements,

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

then for any $t \geq 0$ there is an N so that for $n \geq N$ the group $\langle\langle f_1^n, \ldots, f_m^n \rangle\rangle_G$ has the following properties:

(i) it is isomorphic to F_∞ and has a free basis consisting of conjugates of the f_i^n; and
(ii) the X-translation length of each non-trivial element is at least t.

The main application in this paper is the case where the f_i are pseudo-Anosov elements of the mapping class group and X is the curve complex.

The first statement of Theorem 1.7 will be proved by applying Theorem 1.6 to the projection complex produced by Proposition 3.1. To prove the second statement, we will need to understand translation lengths in terms of projections. Before giving the proof of the theorem, we introduce two tools.

A distance formula for WPD elements. The following proposition gives a lower bound on distance in a hyperbolic space, in terms of projections to quasi-axis bundles. In the statement, d_α is the distance function defined in § 3.2. In addition, for a constant $M \geq 0$ and $x \in \mathbb{R}$ we define $\{\{x\}\}_M$ to be x if $x \geq M$ and zero otherwise.

Proposition 7.1. *Let X be a hyperbolic metric space, let G be a group acting on X, and let $\{f_1, \ldots, f_m\} \subseteq G$ be a collection of WPD elements. Let $\mathcal{A}$ be the collection of quasi-axis bundles in X for the G-conjugates of the f_i. Then there exists a constant M such that for any $x, y \in X$, we have*

$$d(x, y) \geq \frac{1}{6} \sum_{\alpha \in \mathcal{A}} \{\{d_\alpha(x, y)\}\}_M .$$

The formula is reminiscent of the Masur–Minsky distance formula for elements in the mapping class group in terms of subsurface projections [MM00]. This inequality is surely well-known to experts (cf. [BBF15, Theorem 4.13]). We provide a proof in a forthcoming paper.

A lower bound on projections to pivot points. In order to apply Proposition 7.1, we need the following lemma, which says that the local spinning around a pivot point, $d_{v_k}(v_0, h_k v_0)$, is a lower bound on global projection distance, $d_w(v_0, hv_0)$. In the statement, $L(\mathcal{P})$ is the constant from Theorem 1.6.

Lemma 7.2. *Let $\mathcal{P}$ be a projection complex, G a group that acts on $\mathcal{P}$, and $\{R_v\}$ an equivariant L-spinning family subgroups with $L \geq L(\mathcal{P})$. Choose windmill data and let F be the associated free product. Let $h \in F$ and let $h_1 \cdots h_\ell$ be the syllable decomposition of h. Let $1 \leq k \leq \ell$ and let $w_k = h_1 \cdots h_{k-1} v_k \in \mathrm{Piv}(h)$. Then*

$$d_{w_k}(v_0, hv_0) \geq \frac{1}{3} d_{v_k}(v_0, h_k v_0).$$

Proof. By the inequality (1) from the proof of Theorem 1.6, we have that

$$d_{w_k}(v_0, hv_0) \geq d_{v_k}(v_0, h_k v_0) - 2L/3.$$

As $d_{v_k}(v_0, h_k v_0) \geq L$, the proposition follows. □

Proof of Theorem 1.7. We begin with the first statement. For each f_i we fix a quasi-axis bundle $\beta_i \subseteq X$. Let $\mathbb{Y}$ be the set of G-translates of these subsets of X. By Proposition 3.1, the set $\mathbb{Y}$ and the distance functions $\{d_\beta\}_{\beta \in \mathbb{Y}}$ satisfy the projection complex axioms; let $\mathcal{P}$ be this

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

projection complex. Let $L = L(\mathcal{P})$ be the constant from Theorem 1.6 and let τ be the minimum translation length of the f_i. Then, as long as $N \geq L/\tau$, the first statement is an application of Theorem 1.6. It remains to prove the second statement.

Towards proving the second statement, let $\Delta \geq 0$ be such that $\operatorname{diam} \pi_\beta(\beta') \leq \Delta$ for all distinct $\beta, \beta' \in \mathbb{Y}$. (The existence of Δ is implicit in Proposition 3.1; a direct reference for the existence also appears in the paper by Dahmani–Guirardel–Osin [DGO17, Lemma 4.46].) In addition, let M be the constant from Proposition 7.1 applied to X and the collection of quasi-axis bundles $\mathbb{Y}$. We show that the second statement, hence the theorem, holds for $N = (1/\tau)\max\{L, 24\Delta, 4M, 24t\}$. To this end, we let $n \geq N$.

Choose windmill data $\{H_i, W_i, \mathcal{O}_i\}$ and say $W_0 = \{v_0\}$. Suppose $h \in \langle\!\langle f_1^n, \dots, f_m^n \rangle\!\rangle_G$ is non-trivial. Let $h_1 \cdots h_\ell$ be the syllable decomposition of h with respect to the free product data associated to the given windmill data. As translation length is a conjugacy invariant, we may assume that h is cyclically reduced. We treat, in turn, the cases where the syllable length of h is one and where it is greater than one.

If h has syllable length one, then h is a power of some f_i^n, and so for any $x \in X$ we have $d(x, h^p x) \geq pn\tau$ for all $p \geq 0$. Thus, as $n\tau \geq t$ the translation length of h is at least t, as desired.

Now assume that the syllable length of h is at least two. Let $x \in X$ be a point in the quasi-axis bundle $v_0 \in \mathbb{Y}$. To bound $d(x, hx)$ from below, we apply Proposition 7.1, focusing only on the summands corresponding to the quasi-axis bundles in X that are pivot points for h. In order to estimate these terms from below, we apply Lemma 7.2.

Let $w_k = h_1 \cdots h_{k-1} v_k$ be an element of $\operatorname{Piv}(h)$. As h_k is a power of some f_i^n, it follows from Lemma 7.2 that

$$d_{w_k}(v_0, hv_0) \geq \frac{1}{3} d_{v_k}(v_0, h_k v_0) \geq \frac{n\tau}{3}.$$

Hence, as the diameters of $\pi_{w_k}(v_0)$ and $\pi_{w_k}(hv_0)$ are at most Δ, we have

$$d_{w_k}(x, hx) \geq d_{w_k}(v_0, hv_0) - 2\Delta \geq \frac{n\tau}{3} - \frac{n\tau}{12} = \frac{n\tau}{4}.$$

Hence, by Proposition 7.1, the fact that $n\tau/4 \geq M$, and the fact that $n\tau/24 \geq t$, we find

$$d(x, hx) \geq \frac{1}{6} \sum_{w \in \operatorname{Piv}(h)} \{\{d_w(x, hx)\}\}_M \geq \frac{1}{6} \sum_{w \in \operatorname{Piv}(h)} \frac{n\tau}{4} \geq |\operatorname{Piv}(h)|\, t.$$

As h is cyclically reduced and has at least two syllables, we have that $|\operatorname{Piv}(h^p)| = p\,|\operatorname{Piv}(h)|$ for $p \geq 0$ and, hence, the previous argument applied to h^p shows that $d(x, h^p x) \geq p\,|\operatorname{Piv}(h)|\, t$. Thus, the translation length of h in X is at least t, as desired. □

8. Proof of the main theorem

In this section, we prove Theorem 1.1, which states that if G is a subgroup of $\operatorname{Mod}(S)$, $\mathcal{X}$ is a G-overlapping family of subsurfaces, $\mathcal{F}$ is a finite, G-independent family of mapping classes that are carried by $\mathcal{X}$ and are NEC in G, then there is an $N > 0$ with the following properties:

(i) for each $n \geq N$ and any set of orbit representatives $\mathcal{Y}$ for the action of $\langle\!\langle \mathcal{F}^{(n)} \rangle\!\rangle_G$ on $G \cdot \mathcal{X}$ we have

$$\langle\!\langle \mathcal{F}^{(n)} \rangle\!\rangle_G \cong \mathop{\ast}_{Y \in \mathcal{Y}} R_Y;$$

(ii) further, for each $Y \in \mathcal{Y}$ we have $R_Y \cong F_\infty^{b_Y} \times \mathbb{Z}^{a_Y}$.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Symmetrization. Let G, $\mathcal{X}$, and $\mathcal{F}$ be as previously. We explain here how to derive a family of mapping classes from $\mathcal{F}$ that is symmetric in the sense that the subsets of elements supported on two components of $X \in \mathcal{X}$ in the same $\mathrm{Stab}_G(X)$-orbit are conjugate in $\mathrm{Stab}_G(X)$.

First we require some notation. Fix some $X \in \mathcal{X}$. Given a component $X' \in \pi_0(X)$, let $\mathcal{F}_{X'}$ be the subset of $\sigma^{-1}(X)$ consisting of all elements whose support is X'. The set $\sigma^{-1}(X)$ is the disjoint union of $\mathcal{F}_{X'}$ over the components $X' \in \pi_0(X)$.

Suppose now that $X_0, \dots, X_k \in \pi_0(X)$ lie in the same $\mathrm{Stab}_G(X)$-orbit. Let $g_0 \in G$ be the identity and fix elements $g_1, \dots, g_k \in \mathrm{Stab}_G(X)$ such that $g_i X_0 = X_i$. Then consider the set

$$\bigcup_{i,j} (g_j g_i^{-1}) \mathcal{F}_{X_i} (g_j g_i^{-1})^{-1}$$

and let $\widehat{\mathcal{F}}$ be the union of these sets over all $X \in \mathcal{X}$ and all $\mathrm{Stab}_G(X)$-orbits of components of X. We refer to $\widehat{\mathcal{F}}$ as the *symmetrization* of $\mathcal{F}$.

There is an induced function $\hat{\sigma}$ from $\widehat{\mathcal{F}}$ to the set of components of elements of $\mathcal{X}$; this function takes an element of $\widehat{\mathcal{F}}$ to its support.

For $X \in \mathcal{X}$ the normal closure in $\mathrm{Stab}_G(X)$ of the union of the sets $\hat{\sigma}^{-1}(X')$ with X' a component of X is equal to the group R_X, because the new elements are conjugates of the originals by elements in $\mathrm{Stab}_G(X)$. The sets $\hat{\sigma}^{-1}(X')$ are G-independent and each element of $\widehat{\mathcal{F}}$ is NEC.

Applying Theorem 1.7 when the surface has boundary. To understand the group structure of R_X for $X \in \mathcal{X}$, we would like to apply Theorem 1.7 to the action of $\hat{\sigma}^{-1}(X') \subset \mathrm{Stab}_G(X')$ on $\mathcal{C}(X')$ for each non-annular component X' of X. Unfortunately the elements in $\hat{\sigma}^{-1}(X')$ are not necessarily WPD elements for this action. Indeed, if c is a component of $\partial X'$ then T_c acts trivially on $\mathcal{C}(X')$ and so if a power of T_c lies in G then no element of $\mathrm{Stab}_G(X')$ is WPD.

Instead, we proceed as in the discussion on partial pseudo-Anosov mapping classes from the introduction. There is a homomorphism $\mathrm{Stab}_G(X') \to \mathrm{Mod}(\bar{X}')$ obtained by collapsing each component of the boundary to a marked point; let $\bar{G}$ denote the image. As in the proof of Theorem 1.7, there is a natural projection complex that $\bar{G}$ acts on. The vertices correspond to the quasi-axes bundles in $\mathcal{C}(\bar{X}')$ for the images of the elements in $\hat{\sigma}^{-1}(X')$.

Using the action of $\mathrm{Stab}_G(X')$ on this projection complex, we can argue as in the proof of Theorem 1.7 that for any $t \geq 0$ there is an N so that for $n \geq N$ the subgroup

$$\langle\langle f^n \mid f \in \hat{\sigma}^{-1}(X') \rangle\rangle_{\mathrm{Stab}_G(X')}$$

is an infinitely generated free group with basis consisting of conjugates of the various f^n and that the translation length of any non-trivial element in this subgroup on $\mathcal{C}(X')$ is at least t.

Proof of Theorem 1.1. As in the statement, let $\mathcal{X}$ be a G-overlapping family of subsurfaces and let $\mathcal{F} \subset G$ be a finite family of mapping classes that are carried by $\mathcal{X}$ and are G-independent.

If $\mathcal{X} = \{S\}$, then each element of $\mathcal{F}$ is pseudo-Anosov. In particular, $\mathcal{F}$ is a collection of G-independent NEC pseudo-Anosov mapping classes. The theorem then follows from Theorem 1.7, applying the previous discussion in the case that $\partial S \neq \emptyset$.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

We may thus assume in the remainder that $\mathcal{X} \neq \{S\}$. Let $\mathbb{Y} = G \cdot \mathcal{X}$ and let $\mathcal{P}$ be the projection complex obtained by Proposition 3.2. Recall that the distance functions are

$$\underline{d}_Y(X, Z) = \sum_{Y' \in \pi_0(Y)} d_{Y'}(X, Z),$$

where $d_{Y'}$ is the subsurface projection distance. Let $L = L(\mathcal{P})$ be the constant from Theorem 1.6.

Let $\widehat{\mathcal{F}}$ be the symmetrization of $\mathcal{F}$ with corresponding function $\hat{\sigma}$, as defined at the start of the section. Fix some $X \in \mathcal{X}$. For $X' \in \pi_0(X)$ we define

$$\widehat{R}_{X'} = \langle\langle f^n \mid f \in \hat{\sigma}^{-1}(X') \rangle\rangle_{\mathrm{Stab}_G(X')}.$$

The proof proceeds in the following three steps.

(i) There exists an N_X so that for $X' \in \pi_0(X)$ and $n \geq N_X$ the following statements hold:
  (a) $\widehat{R}_{X'}$ is $\mathbb{Z}$ if X' is an annulus and F_∞ otherwise; and
  (b) the translation length of each element of $\widehat{R}_{X'}$ acting on $\mathcal{C}(X')$ is at least L.

(ii) We have $R_X \cong F_\infty^{b_X} \times \mathbb{Z}^{a_X}$.

(iii) The action of R_X on $\mathcal{P}$ is L-spinning.

As $\mathcal{X}$ is finite, the theorem follows from these statements and Theorem 1.6, taking N to be the maximum of the N_X.

Step 1. We define a constant $N_{X'}$ for each component $X' \in \pi_0(X)$ and take N_X to be maximum of these.

If X' is an annulus, then $\hat{\sigma}^{-1}(X')$ consists of a single element f. This element f is a power of the Dehn twist supported on this annulus. We have $d_{X'}(x, f^n x) \geq |n| + 2$ for any $x \in \mathcal{C}(S)$ that intersects X'. As $\langle f \rangle$ is normal in $\mathrm{Stab}_G(X')$ we have that $\widehat{R}_{X'} = \langle f^n \rangle \cong \mathbb{Z}$. In this case, we define $N_{X'}$ to be L.

It remains to treat the case where X' is non-annular. As explained at the start of the section, the proof of Theorem 1.7 shows that there is a constant $N_{X'}$ such that for $n \geq N_{X'}$ the subgroup $\widehat{R}_{X'}$ is an infinitely generated free group with basis consisting of conjugates of the various f^n and such that the translation length of any non-trivial element of $\widehat{R}_{X'}$ on $\mathcal{C}(X')$ is at least L.

Step 2. As X is G-overlapping, it follows that for each $X' \in \pi_0(X)$ the group $\mathrm{Stab}_G(X')$ is a subgroup of $\mathrm{Stab}_G(X)$. There is a natural function

$$\Psi \colon \prod_{X' \in \pi_0(X)} \widehat{R}_{X'} \to \mathrm{Stab}_G(X)$$

that multiplies the coordinate entries of an element in the abstract direct product. The function Ψ is a well-defined homomorphism because elements in the factor subgroups of $\prod_{X' \in \pi_0(X)} \widehat{R}_{X'}$ have support in a single component of X and elements with disjoint support in $\mathrm{Mod}(S)$ commute.

By Step 1(a), we have that $\prod_{X' \in \pi_0(X)} \widehat{R}_{X'} \cong F_\infty^{b_X} \times \mathbb{Z}^{a_X}$. It follows from the definition of symmetrization that the image of Ψ is R_X. It remains to show that Ψ is injective.

Suppose $\Psi(f)$ is trivial and that f is non-trivial. It follows that each coordinate $f_{X'}$ of f is a product of powers of Dehn twists about components of $\partial X'$. Further, there must be at least one non-annular component $X' \in \pi_0(X)$ where $f_{X'}$ is non-trivial, as X contains no parallel annuli. As each non-trivial element of $\widehat{R}_{X'}$ has positive translation length on $\mathcal{C}(X')$ by Step 1(b), and because Dehn twists about components of $\partial X'$ act trivially on $\mathcal{C}(X')$, this is a contradiction.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Step 3. Suppose that $f \in R_X$ is non-trivial and $Y \in \mathbb{Y} - \{X\}$. We write f as a product of elements $f_{X'}$ where $f_{X'} \in \widehat{R}_{X'}$. Then we have

$$\underline{d}_X(Y, fY) = \sum_{X' \in \pi_0(X)} d_{X'}(Y, f_{X'}Y) \geq L,$$

because at least one of the $f_{X'}$ is non-trivial and, hence, by Step 1 its translation length in $\mathcal{C}(X')$ is at least L. This completes the proof. □

9. Applications

In this section, we use Theorem 1.1 to prove two theorems. First we prove Theorem 1.2, which gives an explicit construction of a normal subgroup of $\mathrm{Mod}(S_g)$ that is not contained in any proper level m congruence subgroup. We then prove Theorem 1.5, which gives an explicit example of a pseudo-Anosov mapping class $f \in \mathrm{Mod}(S_g)$ with the property that all non-zero even powers of f normally generate a free subgroup of infinite rank and all odd powers of f normally generate $\mathrm{Mod}(S_g)$.

9.1 Thurston's construction

Our examples will be produced from the Thurston construction of pseudo-Anosov mapping classes. We begin by recalling a special case of Thurston's construction, which works for Dehn twists, and then explain the general case.

Let c and d be curves in S_g that lie in minimal position and have positive geometric intersection number $i(c, d)$. Let X be the surface with marked points obtained as follows: we take a closed regular neighborhood of $c \cup d$ in S_g, take the union of this neighborhood with any complementary regions in S_g that are disks, and then collapse each remaining component of the boundary to a marked point. By construction, the curves c and d fill X, meaning that the complementary regions are disks with at most one marked point. If c and d fill S_g to begin with, then $X = S_g$.

The curves c and d induce a cell decomposition of X (the vertices are the points of $c \cap d$) and the dual complex is a square complex. There is a singular Euclidean structure on X where each 2-cell of the square complex is a Euclidean square. In what follows we take X to be endowed with this structure.

There are two transverse measured foliations on X where the leaves are geodesics parallel to c and d, respectively, and where the transverse measures are given by the metric. We refer to these foliations as the horizontal and vertical foliations of X. These foliations have singularities at the vertices of the square complex.

Let $\mathrm{Aff}(X)$ denote the group of orientation-preserving homeomorphisms of X that preserve the affine structure on X induced by the singular Euclidean metric associated to X. If the complement of the marked points in X has negative Euler characteristic, then homotopic elements of $\mathrm{Aff}(X)$ are equal, and so we may regard $\mathrm{Aff}(X)$ as a subgroup of $\mathrm{Mod}(X)$. We denote by $\mathrm{Isom}(X)$ the finite subgroup of $\mathrm{Aff}(X)$ consisting of isometries of the singular Euclidean metric [FLP91, Expose 9].

The horizontal and vertical foliations on X give an orthonormal frame field, well defined up to sign: at each point, we take unit tangent vectors pointing in the horizontal and vertical directions so that the two vectors (in that order) agree with some fixed orientation on X. With these

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

coordinates we obtain a derivative map

$$D\colon\ \mathrm{Aff}(X) \to \mathrm{PSL}_2\,\mathbb{R}.$$

The Dehn twists T_c and T_d both lie in $\mathrm{Aff}(X)$. Taking $n = i(c,d)$ we have

$$DT_c = \begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad DT_d = \begin{pmatrix} 1 & 0 \\ -n & 1 \end{pmatrix}.$$

Thurston [Thu88, Theorem 7] proved that $f \in \langle T_c, T_d \rangle$ is periodic, reducible, or pseudo-Anosov exactly according to whether Df is elliptic, parabolic, or hyperbolic. In the pseudo-Anosov case, the stretch factor of f is equal to the absolute value of the leading eigenvalue of Df.

The singular Euclidean metric on X associated to the stable and unstable foliations of a pseudo-Anosov $f \in \langle T_c, T_d \rangle$ is equal to the original one defined in terms of the curves c and d. In particular, the stable and unstable foliations are geodesic, and we refer to the corresponding directions in the tangent spaces of the non-singular points as the stable and unstable directions. As the two singular Euclidean structures coincide, the elementary closure of f is a subgroup of $\mathrm{Aff}(X)$.

A multicurve in S_g is a collection of pairwise disjoint curves in S_g. Given a multicurve A, the associated multi-twist is the product of the Dehn twists about the curves in A. Given two multicurves A and B in S_g, there is an analogous Thurston construction. We may form the surface X as in the case where A and B are curves. Instead of a square decomposition we use a rectangle decomposition; the lengths and widths of the rectangles are completely determined by the pairwise intersection numbers of the curves in A and B. As before, there is an associated flat structure on X and a homomorphism $D\colon\ \mathrm{Aff}(X) \to \mathrm{PSL}_2\,\mathbb{R}$. The Nielsen–Thurston type of f is determined by Df in the same way as before.

9.2 Two lemmas

We now state and prove the two lemmas (and a corollary) that will be used to prove Theorems 1.2 and 1.5. The proof of the first lemma is due to Marissa Loving [Lov19].

Lemma 9.1. *Let $g \geq 2$. Let c and d be curves in S_g with $i(c,d) > 0$. Then the mapping classes $f_1 = T_c T_d^{-1}$ and $f_2 = T_c T_d^{-2}$ are normally independent (partial) pseudo-Anosov mapping classes.*

Proof. Let X be the singular Euclidean surface with marked points obtained from c and d as previously and denote $i(c,d)$ by n. We regard f_1 and f_2 as elements of $\mathrm{Mod}(X)$. The mapping classes f_1 and f_2 fall under the Thurston construction. Using the derivative map as before, we see that f_1 and f_2 are pseudo-Anosov with stretch factors

$$\lambda_1 = \frac{n^2 + 2 + n\sqrt{n^2+4}}{2} \quad \text{and} \quad \lambda_2 = n^2 + 1 + n\sqrt{n^2+2}.$$

In particular, λ_1 and λ_2 lie in the quadratic fields $\mathbb{Q}(\sqrt{m+2})$ and $\mathbb{Q}(\sqrt{m})$, where $m = n^2 + 2$.

We claim that $\mathbb{Q}(\sqrt{m+2}) \neq \mathbb{Q}(\sqrt{m})$, which is the same as saying that m and $m+2$ have different square-free parts. If m is even, then $m+2$ is also even, but exactly one of m and $m+2$ is divisible by four, and the claim follows. If m is odd, then no prime factor of m also divides $m+2$, and so again the claim follows.

For $p \in \mathbb{Z}$, the stretch factor of any conjugate of f_i^p is $\lambda_i^{|p|}$. The stretch factor of a (partial) pseudo-Anosov mapping class is irrational, and so λ_i^p is irrational. If a conjugate of f_1^p

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

were equal to a power of f_2 then $\lambda_1^{|p|}$ would be an irrational element of both $\mathbb{Q}(\sqrt{m+2})$ and $\mathbb{Q}(\sqrt{m})$. However, two quadratic fields with a common irrational element are equal. The lemma follows. □

We require one further definition for the statement of the next lemma. Let f be a pseudo-Anosov mapping class with fixed points F_+ and F_- in $\mathrm{PMF}(S_g)$. Recall that the group $\mathrm{EC}(f)$ is the stabilizer of $\{F_+, F_-\}$ in $\mathrm{PMF}(S_g)$, and $\mathrm{EC}^*(f)$ denotes the subgroup fixing both F_+ and F_-. The index of $\mathrm{EC}^*(f)$ in $\mathrm{EC}(f)$ is at most two.

The following lemma was communicated to us by Chris Leininger [Lei19].

Lemma 9.2. *Let A and B be multicurves that fill S_g, and let X be the associated singular Euclidean surface. Let f be a pseudo-Anosov element of $\langle T_A, T_B\rangle$, and let h be a periodic element of $\mathrm{EC}(f)$. Then h preserves $A \cup B$. Moreover, h lies in $\mathrm{EC}^*(f)$ if and only if h preserves both A and B.*

Proof. As h is periodic, Dh is a rotation. Moreover, this rotation must preserve the pair of eigenspaces for Df. If Dh preserves the two eigenspaces, then Dh is trivial, in which case h fixes the horizontal and vertical directions, hence A and B. If Dh interchanges the eigenspaces, then it must be that the stable and unstable directions are orthogonal and that Dh is rotation by $\pi/2$. Thus, Dh also interchanges the horizontal and vertical directions in S_g. The curves of A and B are exactly the curves with vertical or horizontal trajectories, so A and B are interchanged. □

Lemma 9.3. *Let A and B be multicurves in S_g that fill S_g, and assume that there is no element of $\mathrm{Mod}(S_g)$ interchanging A and B. Let f be a pseudo-Anosov element of $\langle T_A, T_B\rangle$. Then f is central in $\mathrm{EC}(f)$. In particular, f is NEC.*

Proof. We first claim that $\mathrm{EC}(f) = \mathrm{EC}^*(f)$. Let $h \in \mathrm{EC}(f)$. If h has infinite order, then h^2 is of infinite order and preserves the points of $\mathrm{PMF}(S_g)$ corresponding to the stable and unstable foliations for f. It follows that h^2 is pseudo-Anosov [FLP91, Exposé 9, Lemme 15]. By the Nielsen–Thurston classification theorem, h is also pseudo-Anosov. As a pseudo-Anosov mapping class has exactly two fixed points in $\mathrm{PMF}(S_g)$, it follows that h^2 and h have the same pair of fixed points, hence $h \in \mathrm{EC}^*(f)$. If h is of finite order, then it follows from Lemma 9.2 and the assumption on A and B that $h \in \mathrm{EC}^*(f)$. This completes the proof of the claim.

Let f_0 be a root of f with minimal stretch factor. There is an internal semidirect product decomposition as follows:

$$\mathrm{EC}^*(f) \cong \langle f_0\rangle \ltimes \mathrm{Isom}(X).$$

The proof of this statement can be found in the unpublished paper by McCarthy [McC85]. Clearly f_0 commutes with f. In addition, because each element of $\mathrm{Isom}(X)$ is of finite order, it follows from Lemma 9.2 and the assumption on A and B that each element of $\mathrm{Isom}(X)$ preserves both A and B, and hence commutes with f. The lemma follows. □

9.3 Proofs of the theorems

We are ready now to prove Theorems 1.2 and 1.5. The construction in the proof of Theorem 1.2 was suggested by Mladen Bestvina [Bes18].

Proof of Theorem 1.2. Choose a non-separating curve c and a separating curve d so that c and d fill the surface S_g. Let $f_1 = T_c T_d^{-1}$ and let $f_2 = T_c T_d^{-2}$. By the Thurston construction, f_1 and f_2

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

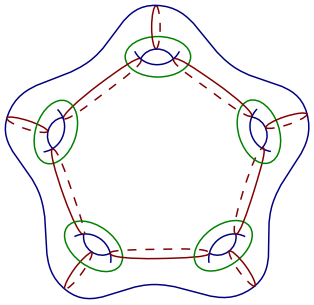

Figure 8. (Colour online) The two multicurves used in the proof of Theorem 1.5.

are pseudo-Anosov, by Lemma 9.1 they are normally independent, and by Lemma 9.3 they are both NEC. We may thus apply Theorem 1.7. Specifically, we may choose distinct prime numbers p_1 and p_2 so that the normal closure N of $f_1^{p_1}$ and $f_2^{p_2}$ is a free group. In particular, N is a proper subgroup of $\mathrm{Mod}(S_g)$.

It remains to show that N is not contained in any congruence subgroup of $\mathrm{Mod}(S_g)$. As T_d acts trivially on $H_1(S_g;\mathbb{Z})$ the action of $f_1^{p_1}$ on $H_1(S_g;\mathbb{Z})$ is the same as that of $T_c^{p_1}$. In particular, $f_1^{p_1}$ lies in $\mathrm{Mod}(S_g)[m]$ if and only if m divides p_1. Similarly, $f_2^{p_2}$ lies in $\mathrm{Mod}(S_g)[m]$ if and only if m divides p_2. Thus, there is no proper subgroup $\mathrm{Mod}(S_g)[m]$ containing N, as desired. □

We were informed by Ashot Minasyan [Min19] of an alternative example a subgroup of $\mathrm{Mod}(S_g)$ satisfying the conclusion of Theorem 1.2. This example is, in fact, not contained in any proper normal subgroup of finite index. The construction uses a result of Michael Hull, and proceeds as follows. Let A be a finitely generated group that has no finite quotients other than the trivial group and set $G = A * A$. Note that G also does not have any finite quotients other than the trivial group. The action of G on the corresponding Bass–Serre tree shows that G is acylindrically hyperbolic. As $\mathrm{Mod}(S_g)$ is also acylindrically hyperbolic [Bow08], it follows from a result of Hull that there is a group Q that is a quotient of both G and $\mathrm{Mod}(S)$ (see [Hul16, Corollary 1.6]). As Q is a quotient of G, it too does not have any finite quotients other than the trivial group. Let K be the kernel of the map $\mathrm{Mod}(S_g) \to Q$. If K is contained in a proper normal subgroup H of finite index in $\mathrm{Mod}(S_g)$, then the image of H in Q is a proper normal finite index subgroup of Q, which is a contradiction as Q does not have any finite quotients other than the trivial group.

The proof of Theorem 1.5 is inspired by the work of the third author with Lanier [LM20, Theorem 1.4]. Specifically, they gave a recipe for constructing a pseudo-Anosov mapping class with the property that all of its odd powers are normal generators for $\mathrm{Mod}(S_g)$. Our construction here is an explicit special case of their recipe, designed so that its even powers have the desired property.

Proof of Theorem 1.5. Fix $g \geq 3$. Let A and B be the multicurves in S_g indicated in Figure 8 (there is only one way to partition the set of curves in the figure into two multicurves). By the Thurston construction the mapping class $f = T_A T_B^{-1}$ is pseudo-Anosov.

Let D_{2g} denote the dihedral group of order $2g$. There is a standard action of D_{2g} on S_g by orientation-preserving homeomorphisms, giving rise to a subgroup of $\mathrm{Mod}(S_g)$ isomorphic to D_{2g}. We refer to this subgroup as simply D_{2g}. Since this action preserves A and B, it follows that D_{2g} lies in the normalizer (indeed, centralizer) of $\langle f \rangle$. In particular, D_{2g} lies in $\mathrm{EC}(f)$.

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press

Let k be an element of D_{2g} corresponding to a reflection of a $2g$-gon and let $h = kf$. We will show that some power of h satisfies the conclusion of the theorem.

First, we claim that all odd powers of h have normal closure equal to $\mathrm{Mod}(S_g)$. Thus, let n be odd. We follow here the argument of Lanier and the third author [LM20, proof of Theorem 1.4]. As k has order two and because k commutes with f, we have that $h^n = kf^n$. Let r denote one of the generators for the cyclic subgroup of D_{2g} of order g. The commutator $[r, h^n] = (rh^n r^{-1})h^{-n}$ lies in the normal closure of h^n. As D_{2g} lies in the centralizer of f we have that

$$[r, h^n] = rh^n r^{-1} h^{-n} = rkf^n r^{-1} kf^{-n} = rkr^{-1}k = r^2,$$

where the last equality uses the relation $kr^{-1}k = r$ in D_{2g}. Lanier and the third author [LM20, Theorem 1.1] showed that for $g \geq 3$ the normal closure of any non-trivial periodic element of $\mathrm{Mod}(S_g)$ besides a hyperelliptic involution is $\mathrm{Mod}(S_g)$. As r^2 is non-trivial and is not a hyperelliptic involution (consider, for instance, the action on $H_1(S_g; \mathbb{Z})$), it follows that the normal closure of r^2, hence h^n, is $\mathrm{Mod}(S_g)$, as desired.

Next, we claim that all sufficiently large even powers of h have normal closure isomorphic to F_∞. If n is even, then $h^n = (f^2)^n$; so a large even power of h is a large power of f. There is no element of $\mathrm{Mod}(S_g)$ interchanging A and B, because these two sets contain different numbers of curves. Thus, by Lemma 9.3, the mapping class f is NEC. By our Theorem 1.7, all sufficiently large powers of f, hence all sufficiently large even powers of h, have normal closure isomorphic to F_∞, as desired.

If n is odd and sufficiently large, then it follows from the previous two claims that h^n satisfies the conclusion of the theorem. □

Acknowledgements

We would like to thank M. Bestvina, K. Bromberg, R. Coulon, F. Dahmani, K. Fujiwara, V. Guirardel, M. Loving, C. Leininger, A. Minasyan, K. Rafi, A. Sisto, and an anonymous referee for helpful comments and conversations. We thank the Mathematical Sciences Research Institute for hosting two of the authors during its 2016 program on Geometric Group Theory, where key parts of this work were completed. The first author is partially supported by the Simons Foundation Grant No. 316383. The second author is supported by National Science Foundation Grant No. DMS-1812021. The third author is supported by National Science Foundation Grants No. DMS-1057874 and No. DMS-1811941.

References

Beh06 J. A. Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space*, Geom. Topol. **10** (2006), 1523–1578; MR 2255505.

Bes18 M. Bestvina, *Personal communication*, May 2018.

BBF13 M. Bestvina, K. Bromberg and K. Fujiwara, *Bounded cohomology via quasi-trees*, Preprint, (2013), arXiv:1306.1542v1.

BBF15 M. Bestvina, K. Bromberg and K. Fujiwara, *Constructing group actions on quasi-trees and applications to mapping class groups*, Publ. Math. Inst. Hautes Études Sci. **122** (2015), 1–64; MR 3415065.

BBFS19 M. Bestvina, K. Bromberg, K. Fujiwara and A. Sisto, *Acylindrical actions on projection complexes*, Enseign. Math. **65** (2019), 1–32; MR 4057354.

BDDKPS20 M. Bestvina, R. Dickmann, G. Domat, S. Kwak, P. Patel and E. Stark, *Free products from spinning and rotating families*, Preprint (2020), arXiv:2010.10735.

BF14 M. Bestvina and M. Feighn, *Hyperbolicity of the complex of free factors*, Adv. Math. **256** (2014), 104–155; MR 3177291.

BF02 M. Bestvina and K. Fujiwara, *Bounded cohomology of subgroups of mapping class groups*, Geom. Topol. **6** (2002), 69–89; MR 1914565 (2003f:57003).

BLM83 J. S. Birman, A. Lubotzky and J. McCarthy, *Abelian and solvable subgroups of the mapping class groups*, Duke Math. J. **50** (1983), 1107–1120; MR 726319.

Bow08 B. Bowditch, *Tight geodesics in the curve complex*, Invent. Math. **171** (2008), 218–300; MR 2367021.

BM19 T. Brendle and D. Margalit, *Normal subgroups of mapping class groups and the metaconjecture of Ivanov*, J. Amer. Math. Soc. **32** (2019), 1009–1070; MR 4013739.

CLM12 M. T. Clay, C. J. Leininger and J. Mangahas, *The geometry of right-angled Artin subgroups of mapping class groups*, Groups Geom. Dyn. **6** (2012), 249–278; MR 2914860.

CLM14 M. Clay, C. J. Leininger and D. Margalit, *Abstract commensurators of right-angled Artin groups and mapping class groups*, Math. Res. Lett. **21** (2014), 461–467; MR 3272023.

CM M. Clay and J. Mangahas, *Hyperbolic quotients of projection complexes*, Groups Geom. Dyn., to appear.

CF J. Crisp and B. Farb, *The prevalence of surface subgroups in mapping class groups*, Preprint.

CP01 J. Crisp and L. Paris, *The solution to a conjecture of Tits on the subgroup generated by the squares of the generators of an Artin group*, Invent. Math. **145** (2001), 19–36; MR 1839284.

CW07 J. Crisp and B. Wiest, *Quasi-isometrically embedded subgroups of braid and diffeomorphism groups*, Trans. Amer. Math. Soc. **359** (2007), 5485–5503; MR 2327038.

Dah18 F. Dahmani, *The normal closure of big Dehn twists and plate spinning with rotating families*, Geom. Topol. **22** (2018), 4113–4144; MR 3890772.

DGO17 F. Dahmani, V. Guirardel and D. Osin, *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*, Mem. Amer. Math. Soc. **245** (2017); MR 3589159.

DHS21 F. Dahmani, M. Hagen and A. Sisto, *Dehn filling Dehn twists*, Proc. Roy. Soc. Edinburgh Sect. A **151** (2021), 28–51; MR 4202630.

Far06 B. Farb, *Some problems on mapping class groups and moduli space*, in *Problems on mapping class groups and related topics*, Proceedings of Symposia in Pure Mathematics, vol. 74 (American Mathematical Society, Providence, RI, 2006), 11–55; MR 2264130.

FM12 B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49 (Princeton University Press, Princeton, NJ, 2012); MR 2850125.

FLP91 A. Fathi, F. Laudenbach and V. Poenaru, *Travaux de Thurston sur les surfaces (Séminaire Orsay)* (Société Mathématique de France, Paris, 1991). Reprint of *Travaux de Thurston sur les surfaces*, Astérisque No. 66-67 (Société Mathématique de France, Paris, 1979) [MR 0568308 (82m:57003)]; MR 1134426.

Fuj15 K. Fujiwara, *Subgroups generated by two pseudo-Anosov elements in a mapping class group. II. Uniform bound on exponents*, Trans. Amer. Math. Soc. **367** (2015), 4377–4405; MR 3324932.

Gro01 M. Gromov, CAT(κ)*-spaces: construction and concentration*, Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI) **280** (2001), no. Geom. i Topol. 7, 100–140, 299–300; MR 1879258.

HT02 H. Hamidi-Tehrani, *Groups generated by positive multi-twists and the fake lantern problem*, Algebr. Geom. Topol. **2** (2002), 1155–1178; MR 1943336.

HPW15 S. Hensel, P. Przytycki and R. C. H. Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, J. Eur. Math. Soc. (JEMS) **17** (2015), 755–762; MR 3336835.

Hul16 M. Hull, *Small cancellation in acylindrically hyperbolic groups*, Groups Geom. Dyn. **10** (2016), 1077–1119; MR 3605028

Ish96 A. Ishida, *The structure of subgroup of mapping class groups generated by two Dehn twists*, Proc. Japan Acad. Ser. A Math. Sci. **72** (1996), 240–241; MR 1435728

Iva92 N. V. Ivanov, *Subgroups of Teichmüller modular groups*, Translations of Mathematical Monographs, vol. 115 (American Mathematical Society, Providence, RI, 1992). Translated from the Russian by E. J. F. Primrose and revised by the author. MR 1195787.

Iva06 N. V. Ivanov, *Fifteen problems about the mapping class groups*, in *Problems on mapping class groups and related topics*, Proceedings of Symposia in Pure Mathematics, vol. 74 (American Mathematical Society, Providence, RI, 2006), 71–80; MR 2264532.

Kob12 T. Koberda, *Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups*, Geom. Funct. Anal. **22** (2012), 1541–1590; MR 3000498.

LM18 J. Lanier and D. Margalit, *Normal generators for mapping class groups are abundant (first version)*, Preprint (2018), arXiv:1805.03666v1.

LM20 J. Lanier and D. Margalit, *Normal generators for mapping class groups are abundant*, Preprint (2020), arXiv:1805.03666.

Lei19 C. Leininger, *Personal communication*, January 2019.

Lön10 M. Lönne, *Presentations of subgroups of the braid group generated by powers of band generators*, Topology Appl. **157** (2010), 1127–1135; MR 2607077.

Lov19 M. Loving, *Personal communication*, January 2019.

Man10 J. Mangahas, *Uniform uniform exponential growth of subgroups of the mapping class group*, Geom. Funct. Anal. **19** (2010), 1468–1480; MR 2585580 (2011d:57002).

Man13 J. Mangahas, *A recipe for short-word pseudo-Anosovs*, Amer. J. Math. **135** (2013), 1087–1116; MR 3086070.

MR20 K. Mann and K. Rafi, *Large scale geometry of big mapping class groups*, Preprint (2020), arXiv:1912.10914.

MM99 H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), 103–149; MR 1714338 (2000i:57027).

MM00 H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. II. Hierarchical structure*, Geom. Funct. Anal. **10** (2000), 902–974; MR 1791145 (2001k:57020).

McC85 J. McCarthy, *A "Tits-alternative" for subgroups of surface mapping class groups*, Trans. Amer. Math. Soc. **291** (1985), 583–612; MR 800253.

Min19 A. Minasyan, *Personal communication*, June 2019.

Run20 I. Runnels, *Effective generation of right-angled Artin groups in mapping class groups*, Preprint (2020), arXiv:2004.13585.

Seo19 D. Seo, *Powers of Dehn twists generating right-angled Artin groups*, Preprint (2019), arXiv:1909.03394.

Ser03 J.-P. Serre, *Trees*, Springer Monographs in Mathematics (Springer, Berlin, 2003), translated from the French original by John Stillwell, corrected 2nd printing of the 1980 English translation; MR 1954121 (2003m:20032).

Thu88 W. P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) **19** (1988), 417–431; MR 956596.

Whi00 K. Whittlesey, *Normal all pseudo-Anosov subgroups of mapping class groups*, Geom. Topol. **4** (2000), 293–307; MR 1786168.

RAAGs as normal subgroups of mapping class groups

Matt Clay mattclay@uark.edu
Department of Mathematical Sciences, University of Arkansas, 309 SCEN, Fayetteville, AR 72701, USA

Johanna Mangahas mangahas@buffalo.edu
Department of Mathematics, University at Buffalo, 244 Mathematics Building, Buffalo, NY 14260, USA

Dan Margalit margalit@math.gatech.edu
School of Mathematics, Georgia Institute of Technology, 686 Cherry Street, Atlanta, GA 30332, USA

https://doi.org/10.1112/S0010437X21007417 Published online by Cambridge University Press