

RIGHT-ANGLED ARTIN GROUPS AND A GENERALIZED ISOMORPHISM PROBLEM FOR FINITELY GENERATED SUBGROUPS OF MAPPING CLASS GROUPS

THOMAS KOBERDA

ABSTRACT. Consider the mapping class group $\text{Mod}_{g,p}$ of a surface $\Sigma_{g,p}$ of genus g with p punctures, and a finite collection $\{f_1, \dots, f_k\}$ of mapping classes, each of which is either a Dehn twist about a simple closed curve or a pseudo-Anosov homeomorphism supported on a connected subsurface. In this paper we prove that for all sufficiently large N , the mapping classes $\{f_1^N, \dots, f_k^N\}$ generate a right-angled Artin group. The right-angled Artin group which they generate can be determined from the combinatorial topology of the mapping classes themselves. When $\{f_1, \dots, f_k\}$ are arbitrary mapping classes, we show that sufficiently large powers of these mapping classes generate a group which embeds in a right-angled Artin group in a controlled way. We establish some analogous results for real and complex hyperbolic manifolds. We also discuss the unsolvability of the isomorphism problem for finitely generated subgroups of $\text{Mod}_{g,p}$, and prove that the isomorphism problem for right-angled Artin groups is solvable. We thus characterize the isomorphism type of many naturally occurring subgroups of $\text{Mod}_{g,p}$.

March 21, 2012

CONTENTS

1. Introduction and statement of results	2
1.1. Overview	2
1.2. Right-angled Artin subgroups of mapping class groups	2
1.3. Remarks on embedding right-angled Artin groups in mapping class groups	4
1.4. Commensurability questions	4
1.5. Motivation from linear groups	5
1.6. Decision problems	5
1.7. Analogies with lattices and symmetric spaces	6
1.8. Notes and open questions	7
2. Acknowledgements	8
3. The ping-pong lemma as the fundamental tool	8
4. Analogy with finite volume hyperbolic k -manifolds and rank one Lie groups	9
4.1. Real hyperbolic spaces	9
4.2. Complex hyperbolic spaces	10
4.3. Commensurability and embeddings	11
5. The bad behavior of right-angled Artin groups	14
6. Right-angled Artin groups as subgroups of mapping class groups	17
6.1. Reducible mapping classes	17
6.2. Convergence in $\mathbb{P}\mathcal{ML}(\Sigma)$	19
6.3. Intersections and freeness	19
6.4. Competing convergence	20
6.5. The issue with boundary twisting	21
6.6. Angles of intersection	22

1991 *Mathematics Subject Classification*. Primary 37E30; Secondary 20F36, 05C60.

Key words and phrases. Subgroups of mapping class groups, right-angled Artin groups.

6.7. Fellow traveling	25
6.8. Greedy normal forms in right-angled Artin groups	25
6.9. Playing ping-pong and the proof of Theorem 1.1	26
7. Right-angled Artin groups in $\text{Mod}_{g,p}$ and the isomorphism problem for subgroups of $\text{Mod}_{g,p}$	28
7.1. Decision problems in $\text{Mod}_{g,p}$	28
7.2. Homological rigidity of right-angled Artin groups: no accidental isomorphisms	30
7.3. Embedding right-angled Artin groups in mapping class groups	32
8. Embeddings and cohomology	35
8.1. Commensurability questions	35
8.2. Virtual embedding questions	35
8.3. Virtual speciality questions	37
References	37

1. INTRODUCTION AND STATEMENT OF RESULTS

1.1. Overview. The goal of this paper is to describe a large class of finitely generated subgroups of the mapping class group. Several years ago, B. Farb conjectured that given any finite collection of mapping classes, sufficiently large powers of these classes generate a right-angled Artin group. We resolve Farb’s conjecture in this article.

The task of describing all finitely generated subgroups up to isomorphism is an impossible one, as we will show precisely.

1.2. Right-angled Artin subgroups of mapping class groups. Let $\Sigma = \Sigma_{g,p}$ be a surface of finite genus g and a finite set P of $p \geq 0$ punctures. We will require that

$$\chi(\Sigma) = 2 - 2g - p < 0.$$

Recall that the mapping class group $\text{Mod}_{g,p}$ is defined by

$$\text{Mod}_{g,p} \cong \pi_0(\text{Homeo}(\Sigma, P)),$$

namely the group of orientation-preserving self-homeomorphisms of Σ which preserve the set of punctures, up to isotopy. It is sometimes important to distinguish between mapping classes which preserve P pointwise and those which preserve P setwise. This distinction is not important for our considerations, as the former sits within the latter with finite index.

We first recall the notion of a right-angled Artin system and a right-angled Artin group. Let Γ be a finite graph with vertex set V and edge set E . The set of vertices of this graph together with their incidence (which is to say adjacency) relations determine a **right-angled Artin system**. This terminology is motivated Coxeter theory. A **right-angled Artin group** is a pair $(A(\Gamma), V)$, where V is a right-angled Artin system and a group $A(\Gamma)$ with a presentation given by

$$A(\Gamma) = \langle V \mid [v_i, v_j] \text{ whenever } (v_i, v_j) \in E \rangle.$$

We call the size of the vertex set of Γ the **rank** of $A(\Gamma)$, which is terminology motivated by the fact that this is the smallest number of generators for $A(\Gamma)$. When the right-angled Artin system for $A(\Gamma)$ is the set of vertices of Γ , we call the vertices of Γ **vertex generators**.

For curves and connected subsurfaces of a given surface Σ , there is a notion of a right-angled Artin system as well. Let F be a nonempty, finite collection of Dehn twists or pseudo-Anosov homeomorphisms supported on connected subsurfaces of Σ . The support Σ_i of each element $f_i \in F$ is either a simple closed curve or a connected subsurface. We can build a graph Γ by taking the vertices of Γ to be the elements of F and connecting two vertices by an edge if and only if the

supports can be realized disjointly. This graph is called the **coincidence correspondence of F** . We define $A(F)$ to be $A(\Gamma)$.

Call a subset $F = \{f_1, \dots, f_k\} \subset \text{Mod}_{g,p}$ **irredundant** if there are no linear relations between commuting subcollections of F . Precisely, let $\{f_1, \dots, f_i\} = F' \subset F$ be a collection of elements which generate an abelian subgroup of $\text{Mod}_{g,p}$. The elements F' satisfy a **linear relation** if (after writing the group law in this subgroup additively) we have a nontrivial solution to the equation

$$a_1 f_1 + \dots + a_i f_i = 0.$$

The main result in this paper can be stated as follows:

Theorem 1.1. *Let $\{f_1, \dots, f_k\} \subset \text{Mod}_{g,p}$ be an irredundant collection of Dehn twists or pseudo-Anosov homeomorphisms supported on connected subsurfaces of a given surface. Then there exists an N such that for each $n \geq N$, the set of mapping classes $\{f_1^n, \dots, f_k^n\}$ is a right-angled Artin system for a right-angled Artin subgroup of $\text{Mod}_{g,p}$. The isomorphism type of this right-angled Artin group is given by $A(\Gamma)$, where Γ is the coincidence correspondence of $\{f_1, \dots, f_k\}$.*

We remark that in the hypotheses of Theorem 1.1, irredundancy is equivalent to the requirement that no two mapping classes $\{f_i, f_j\}$ for $i \neq j$ generate a cyclic subgroup of $\text{Mod}_{g,p}$. There is a minor issue which arises in relation to two pseudo-Anosov mapping classes ψ and ϕ supported on a subsurface $\Sigma' \subset \Sigma$ which commute but which do not share a common power. In this case, ψ and ϕ differ by a multitwist about the boundary components of Σ' . We will adopt a convention throughout this paper which disallows such a setup. If we allow ψ on our list of mapping classes then we will insist that ψ does not twist about the boundary components of Σ' . We will consider ϕ as a composition of ψ with a multitwist and hence as a mapping class with disconnected support. However, we will allow the separate twists about the boundary components of Σ' to be on our list. We discuss this issue in more detail in Subsection 6.5.

Theorem 1.1 says that after perhaps passing to powers, there are few non-obvious relations between mapping classes. This means that the various relations in mapping class groups such as braid relations, chain relations, lantern relations, etc. are all “shallow” in the mapping class group.

We will say that a group G is **enveloped by a right-angled Artin group $A(\Gamma)$** if G is generated by elements contained in the **clique subgroups** of Γ , which is to say the subgroups induced by complete subgraphs of Γ .

The following result is an immediate corollary of Theorem 1.1 and the description of reducible mapping classes.

Corollary 1.2. *Let $\{f_1, \dots, f_k\}$ be any collection of mapping classes. Then there exists an N such that for each $n \geq N$, the group $\langle f_1^n, \dots, f_k^n \rangle$ is enveloped by a right-angled Artin subgroup of the mapping class group.*

As we will show below, it is not true that a group enveloped by a right-angled Artin group is itself a right-angled Artin group, even abstractly. Furthermore, we will see that any group enveloped by a right-angled Artin group can be embedded in some mapping class group, so the conclusion of Theorem 1.1 is as strong as one could hope.

To get the full conclusion of Theorem 1.1, it is clear that irredundancy is a necessary hypothesis. It is clear that the result also fails if any one of the mapping classes $\{f_1, \dots, f_k\}$ has finite order. The hypothesis that all the mapping classes be Dehn twists or pseudo-Anosov on subsurfaces is likewise essential, for the reason that not every group enveloped by a right-angled Artin group is itself a right-angled Artin group.

An immediate question raised by Theorem 1.1 asks which right-angled Artin groups occur as subgroups of mapping class groups. For a fixed genus, ascertaining which right-angled Artin groups which embed in $\text{Mod}_g = \text{Mod}_{g,0}$ appears to be a rather subtle problem. However, we can show the following:

Proposition 1.3. *Let $A(\Gamma)$ be a right-angled Artin group. There exists a g such that $A(\Gamma) < \text{Mod}_g$, with the image of the vertices of Γ consisting of powers of Dehn twists about simple closed curves or pseudo-Anosov homeomorphisms supported on connected subsurfaces. It is possible to arrange the embedding in such a way that the image is contained in the Torelli subgroup of Mod_g .*

1.3. Remarks on embedding right-angled Artin groups in mapping class groups. In [15], it was shown that every right-angled Artin group can be embedded in some right-angled Artin group, with a right-angled Artin system given by a collection of powers of Dehn twists about simple closed curves. A result of similar ilk can be found in the paper [16] of Crisp and Wiest. If Γ is a finite graph, its **complement** (or **dual**) graph Γ^* is given by embedding Γ into the complete graph on the vertices of Γ , and then deleting the edges which belong to Γ . Crisp and Wiest prove that if Γ^* embeds in a surface of genus g then the right-angled Artin group $A(\Gamma)$ embeds in a genus g surface braid group. The problem of embedding various right-angled Artin groups into mapping class groups is the subject of Crisp and Farb’s preprint [14]. Sabalka showed in [50] that right-angled Artin groups can always be embedded into graph braid groups.

As an immediate corollary to Proposition 1.3, we obtain:

Corollary 1.4. *Let $A(\Gamma)$ be a right-angled Artin group and let*

$$\langle g_1, \dots, g_k \rangle = G < A(\Gamma)$$

be a finitely generated subgroup. Then there exists an N such that for all $n \geq N$, we have that

$$\langle g_1^n, \dots, g_k^n \rangle$$

is enveloped by a right-angled Artin group.

We record this corollary because it is well-known that there are many extremely complicated subgroups of right-angled Artin groups when the flag complex associated to the right-angled Artin system is not simply connected (see [5], also [19]). A systematic study of right-angled Artin subgroups of a given right-angled Artin group has recently been carried out by S. Kim and the author in [37].

1.4. Commensurability questions. Recall that two groups G and H are **commensurable** if there exist finite index subgroups G' and H' of G and H respectively such that $G' \cong H'$. Since $\text{Mod}_{g,p}$ is finitely generated by Dehn twists, it is therefore natural to ask whether or not $\text{Mod}_{g,p}$ is commensurable with a right-angled Artin group. More generally, one can wonder whether $\text{Mod}_{g,p}$ virtually injects into a right-angled Artin group.

C. Leininger has informed the author that combining the results of Crisp and Wiest in [16] and Leininger and Reid in [42], one can show that $\text{Mod}_{2,0}$ is not commensurable with a right-angled Artin group. A proof of this fact appears in [13]. For higher genera, an even stronger conclusion is true:

Theorem 1.5. *If $g > 2$ or if $g = 2$ and $p > 1$ then $\text{Mod}_{g,p}$ does not virtually inject into a right-angled Artin group.*

We thus recover a corollary to a result of M. Kapovich and B. Leeb (see [35]), namely that a sufficiently complicated ($g > 2$ or $g = 2$ and $p > 1$) mapping class group does not embed into any right-angled Artin group. Kapovich and Leeb prove a slightly different statement, namely that sufficiently complicated mapping class groups (which is to say $g = 2$ and $p \geq 1$ or $g > 2$) do not act effectively, cocompactly and discretely on Hadamard spaces, which is to say Riemannian manifolds with everywhere nonpositive sectional curvatures. In modern terminology, Theorem 1.5 can be stated as “mapping class groups are not virtually special” (cf. the recent work of Haglund and Wise in [30]).

Theorem 1.5 may hold for closed surfaces of genus two, but the methods in this paper do not work when Σ has genus 2, and the author is not aware of a way to rectify this difficulty. Note that

we have $\text{Mod}_{1,0} \cong \text{Mod}_{1,1} \cong SL_2(\mathbb{Z})$, which is virtually a free group and hence commensurable with a right-angled Artin group.

Along lines similar to Theorem 1.5, M. Bridson has shown in [8] that whenever the mapping class group of a closed orientable surface of genus g acts by semisimple isometries on a complete CAT(0) space of dimension less than g , it has a global fixed point. Actions of mapping class groups by isometries on CAT(0) spaces are also studied in [7]. Bridson also shows that a genus 2 mapping class group acts properly by semisimple isometries on an 18-dimensional complete CAT(0) space.

1.5. Motivation from linear groups. One of the motivations for Theorem 1.1 comes from linear groups and the **Tits alternative** (see [55]), which was originally formulated for finitely generated linear groups:

Theorem 1.6. *Let G be a finitely generated linear group. Then G is either virtually solvable or contains a nonabelian free group.*

It is known that mapping class groups satisfy the Tits alternative. One of the main results of Birman Lubotzky and McCarthy (see [6]) is the following:

Theorem 1.7. *Let $G < \text{Mod}_{g,p}$ be virtually solvable. Then G is virtually abelian. Furthermore, the rank of a free abelian subgroup of G is bounded by the size of a maximal collection of disjoint, non-peripheral, pairwise non-isotopic simple closed curves on $\Sigma_{g,p}$.*

A subgroup of the mapping class group which is virtually abelian but not virtually cyclic cannot contain pseudo-Anosov elements. By the work of Ivanov (see [33]), such a subgroup must be reducible, which means that it stabilizes a finite nonempty collection of isotopy classes of essential simple closed curves. Ivanov also proved that every subgroup of $\text{Mod}_{g,p}$ which is irreducible and not virtually cyclic contains a nonabelian free group, completing the Tits alternative for $\text{Mod}_{g,p}$.

With our setup, the Tits alternative can be refined to read as follows:

Theorem 1.8. *Let $G < \text{Mod}_{g,p}$ be generated by $\{f_1, \dots, f_k\}$, where each f_i has infinite order. Then exactly one of the following holds:*

- (1) G is virtually torsion-free abelian.
- (2) There is an N such that for all $n \geq N$, we have that $\{f_1^n, \dots, f_k^n\}$ generates a nonabelian subgroup $H < G$ which is enveloped by a right-angled Artin subgroup of $\text{Mod}_{g,p}$.

Any two-generated nonabelian group enveloped by a right-angled Artin group is automatically free, as can be deduced from Lemma 3.1. Thus, we do indeed obtain a strengthening of the Tits alternative.

1.6. Decision problems. Another perspective on the discussion in this paper stems from decision problems in group theory. In general, mapping class groups are so complicated that there is no algorithm for computing whether or not two finitely generated subgroups are isomorphic. We shall prove a more general fact:

Proposition 1.9. *Let G be a finitely presented group which contains an embedded product of two nonabelian free groups. There is no algorithm which determines whether two finitely generated subgroups of G are isomorphic.*

We show that solving the isomorphism problem for finitely generated subgroups of $\text{Mod}_{g,p}$ solves the generation problem for $F_6 \times F_6$, which is known to be unsolvable by the work of Lyndon and Schupp ([43]) and C. Miller ([46]). We will use this latter fact to show that if a group contains a copy of $F_2 \times F_2$ then it has an unsolvable isomorphism problem for finitely generated subgroups by showing that the isomorphism problem for finitely generated subgroups of $F_2 \times F_2$ solves the generation problem for $F_6 \times F_6$.

On the other hand, the isomorphism type of a right-angled Artin group can be algorithmically determined. To illustrate this fact, we prove the following rigidity result:

Theorem 1.10. *Let $A(\Gamma)$ and $A(\Delta)$ be two right-angled Artin groups. Then $A(\Gamma) \cong A(\Delta)$ if and only if $H^*(A(\Gamma), \mathbb{Q}) \cong H^*(A(\Delta), \mathbb{Q})$ as algebras, which happens if and only if Γ and Δ are isomorphic as graphs.*

Theorem 1.10 can be found in Sabalka’s paper [51], and a proof follows from a combination of theorems of Droms in [19] and Gubeladze in [29].

The fact that $A(\Gamma) \cong A(\Delta)$ if and only if $\Gamma \cong \Delta$ was proved by C. Droms in [20]. Combining Theorem 1.10 with Droms’ result, we see that to specify the isomorphism type of a right-angled Artin group, the rational cohomology algebra and the graph both suffice.

If $\{f_1, \dots, f_n\}$ are mapping classes which are a right-angled Artin system for a right-angled Artin group $A(\Gamma)$, the structure of Γ can be gleaned from the intersection theory of subsurfaces and reduction systems for these mapping classes. It follows that Theorem 1.1 and Theorem 1.10 together provide a partial solution to an “up to powers” isomorphism problem for finitely generated subgroups of the mapping class group, although the isomorphism problem for finitely generated subgroups of the mapping class groups is unsolvable.

1.7. Analogies with lattices and symmetric spaces. The final perspective on Theorem 1.1 which we explore comes from the analogy between Teichmüller spaces and symmetric spaces, and between mapping class groups and lattices in semisimple Lie groups. We have seen that each right-angled Artin group is contained in a mapping class group, so we can ask which right-angled Artin groups are contained in lattices in semisimple Lie groups. Similarly, we can ask whether a statement analogous to Theorem 1.1 holds. Questions of the former ilk have been discussed by T. Januszkiewicz and J. Świątkowski in [34] and by S. Wang in [56], for instance.

It turns out that the conclusion of Theorem 1.1 fails dramatically for lattices in $SL_n(\mathbb{R})$ for instance, when $n \geq 3$. The following result is due to H. Oh in [48]. The precise definitions of all the terms are not important for understanding the spirit of the result:

Theorem 1.11. *Let G be an absolutely adjoint simple real Lie group of \mathbb{R} -rank at least two and let $U_1, U_2 < G$ be opposite horospherical subgroups. Suppose that G is not split over \mathbb{R} and that U_1 is not the unipotent radical of a Borel subgroup in a group of type A_2 .*

Let F_1 and F_2 be lattices in $U_1(\mathbb{R})$ and $U_2(\mathbb{R})$ respectively and let Γ be the subgroup generated by these two lattices. If Γ is discrete then there is a \mathbb{Q} -form with respect to which U_1 and U_2 are defined over \mathbb{Q} and F_i is commensurable with $U_i(\mathbb{Z})$ for each i . Furthermore, Γ is commensurable with $G(\mathbb{Z})$.

The requirement in the previous result that G does not split over \mathbb{R} is essential. For instance, one obtains immediate generalizations of Theorem 1.1 for groups like $SL_2(\mathbb{Z}) \times SL_2(\mathbb{Z})$, since we can view $SL_2(\mathbb{Z}) \times SL_2(\mathbb{Z})$ as the mapping class group of a stable curve consisting of two tori glued together at a point.

To contrast Oh’s result with Theorem 1.1, the role of the unipotent subgroups could be played by powers of Dehn twists about pairwise intersecting curves. Theorem 1.1 implies that powers of those Dehn twists generate a free group, which is far from being the appropriate analogue of a lattice in mapping class group theory. So, Theorem 1.1 can be thought of as a “rank one” phenomenon for mapping class groups.

For discrete subgroups of rank one semisimple Lie groups, we do get an analogue of Theorem 1.1, though the right-angled Artin groups which occur are limited. We have the following results:

Proposition 1.12. *Let $\{g_1, \dots, g_m\} \subset SO(k, 1)$ generate a discrete subgroup, and suppose furthermore that these elements are irredundant. Then there exists an N such that for all $n \geq N$, we have that $\{g_1^n, \dots, g_m^n\}$ is a right-angled Artin system for a free product of finitely generated free abelian subgroups of $SO(k, 1)$.*

For some other rank one Lie groups one gets very close but not verbatim analogues of Theorem 1.1. If \mathcal{N} is a class of groups and G is a free product of \mathcal{N} -groups, we call a subset $S \subset G$ a

right-angled Artin system for G if $\langle S \rangle = G$ and if each $s \in S$ is contained in a free factor for some free product decomposition of G .

Proposition 1.13. *Let $\{g_1, \dots, g_m\} \subset SU(k, 1)$ generate a discrete subgroup, and suppose furthermore that these elements are irredundant. Then there exists an N such that for all $n \geq N$, we have that $\{g_1^n, \dots, g_m^n\}$ is a right-angled Artin system for a free product of finitely generated torsion-free nilpotent subgroups of $\text{Isom}(\mathbb{H}_{\mathbb{C}}^k)$ under the projection*

$$SU(k, 1) \rightarrow \text{Isom}(\mathbb{H}_{\mathbb{C}}^k),$$

where $\mathbb{H}_{\mathbb{C}}^k$ denotes complex hyperbolic space.

By analogy to Theorem 1.5, we will show:

Theorem 1.14. *Let $k \geq 3$ and let $\Gamma < SO(k, 1)$ be a lattice. Then Γ is not commensurable with a right-angled Artin group.*

If one removes the hypothesis of discreteness in any of the results above about rank one Lie groups, the conclusions are false. In $SL_2(\mathbb{R})$, for instance, there are finitely generated (indiscrete) solvable subgroups such that no powers of the generators generate an abelian group. B. Farb has conjectured an analogue of Theorem 1.1 for $SL_2(\mathbb{R})$ which would account for the aforementioned solvable subgroups, but his conjecture appears to be extremely difficult to establish.

1.8. Notes and open questions. Right-angled Artin subgroups of mapping class groups have been studied by various authors, for example by Crisp and Paris in [15]. Our methods exploit the dynamics of the action of $\text{Mod}_{g,p}$ on the space of projective measured laminations on Σ , i.e. $\mathbb{P}\mathcal{ML}(\Sigma)$. An account of some aspects of this theory can be found in the paper of McCarthy and Papadopoulos ([45]). Questions about right-angled Artin groups in mapping class groups generated by powers of mapping classes were considered by L. Funar in [27] and M. Lönne in [41].

A recent result of Clay, Leininger and Mangahas closely related to Theorem 1.1 analyzes right-angled Artin groups which quasi-isometrically embed into mapping class groups and can be found in [13]. The methods of their paper rely heavily on the complex of curves.

An interesting related question is to extend the theory in this paper to $\text{Out}(F_n)$. It is unclear whether or not the exact analogue of Theorem 1.1 holds for $\text{Out}(F_n)$, and the proof given here will certainly not work for $\text{Out}(F_n)$. One can deduce an analogue to Theorem 1.5 for $\text{Out}(F_n)$ when $n \geq 4$ as follows: note that $\text{Aut}(F_n) < \text{Out}(F_{n+1})$. By the result of Formanek and Procesi in [25], $\text{Aut}(F_n)$ is not linear when $n > 2$, so that $\text{Out}(F_n)$ is nonlinear for $n > 3$. It follows that $\text{Out}(F_n)$ cannot virtually inject into a right-angled Artin group for $n > 3$. One can ask:

Question 1.15. *Does $\text{Out}(F_3)$ virtually embed in a right-angled Artin group?*

A positive answer to this question would immediately imply that $\text{Out}(F_3)$ is linear. Furthermore, it would immediately imply all sorts of properties for $\text{Out}(F_3)$, such as virtual residually torsion-free nilpotence and the Tits alternative, and it would furnish $\text{Out}(F_3)$ with a properly discontinuous action by isometries on a CAT(0) space.

Finally, we pose the following question:

Question 1.16. *Let f_1, \dots, f_k be mapping classes. Is there a uniform N (which depends only on the underlying surface) such that $\{f_1^N, \dots, f_k^N\}$ generate a subgroup which is enveloped by a right-angled Artin subgroup of $\text{Mod}_{g,p}$?*

This question appears to be open in general. When $\{\psi, \phi\}$ is a pair of pseudo-Anosov mapping classes, an affirmative answer was given by Fujiwara in [26]. A related statement has been conjectured by Funar in [27].

2. ACKNOWLEDGEMENTS

The author thanks B. Farb for suggesting the problem, for his help and comments, and for sharing the preprint [14]. The author also thanks J. Aramayona, M. Bridson, M. Casals, T. Church, M. Clay, U. Hamenstädt, C. Leininger, J. Mangahas, C. McMullen, C.F. Miller, A. Putman, M. Sapir, A. Silberstein and K. Vogtmann for useful conversations and comments. The author especially thanks M. Kapovich for a careful reading of earlier versions of the manuscript. The author thanks the Hausdorff Research Institute for Mathematics in Bonn and the Institute for Mathematical Sciences in Singapore, where this research was completed. The author finally thanks the referee for very helpful comments which greatly improved the exposition and correctness of the article. The author is partially supported by an NSF Graduate Research Fellowship.

3. THE PING-PONG LEMMA AS THE FUNDAMENTAL TOOL

When a free group acts on a set X , one often wants to know whether or not it acts faithfully. Given any candidates for two free factors, if these two factors “switch” two subsets of X then under some further technical hypotheses, the action will be faithful. A more general ping-pong argument holds for right-angled Artin groups and can be found in the preprint [14] of Crisp and Farb. We recall a proof here for the reader’s convenience. A much more detailed exposition on ping-pong lemmas and the material in Section 4 can be found in [38]. Let Γ be a simplicial graph with vertex set $\{1, \dots, n\}$ and edge set $E(\Gamma)$. We denote the associated right-angled Artin group by $A(\Gamma)$. The group $A(\Gamma)$ is generated by elements $\{g_1, \dots, g_n\}$ according to the standard presentation. We will interchangeably call those generators the **vertex generators** or a **right-angled Artin system**. The ping-pong argument can be stated as follows:

Lemma 3.1. *Let X be a set on which $A(\Gamma)$ acts. Suppose there exist (not necessarily disjoint) subsets $X_1, \dots, X_n \subset X$ and an element*

$$x \in X \setminus \bigcup_{i=1}^n X_i$$

such that for each nonzero integer k , we have:

- (1) $g_i^k(X_j) \subset X_j$ if $(i, j) \in E(\Gamma)$.
- (2) $g_i^k(X_j) \subset X_i$ if $(i, j) \notin E(\Gamma)$.
- (3) $g_i^k(x) \in X_i$ for all i .

Then $A(\Gamma)$ acts faithfully on X .

First proof of Lemma 3.1. Under the hypotheses it will follow that the orbit of x is a transitive $A(\Gamma)$ -set. We write $\ell(g)$ for the word length of $g \in A(\Gamma)$ with respect to the standard generating set. It suffices to show that if $g = g_i^{\pm 1}g'$ with $\ell(g) = \ell(g') + 1$, then $g(x) \in X_i$, for then no non-identity element of $A(\Gamma)$ fixes x .

We proceed by induction on $\ell(g)$. We write $g = g_i^k w$, with $\ell(g) = \ell(w) + |k|$, $k \neq 0$ and $|k|$ maximal. If $w = 1$ then $g(x) = g_i^k(x) \in X_i$. Otherwise, $w = g_j w'$, with $\ell(w) = \ell(w') + 1$ and $i \neq j$. If $(i, j) \notin E(\Gamma)$, we obtain that $w(x) \in X_j$ and $g(x) \in X_i$. If $(i, j) \in E(\Gamma)$ then $g = g_j g_i^k w'$, where $\ell(g) = \ell(g_i^k w') + 1$ so that $g_i^k w'(x) \in X_i$. It then follows that $g(x) \in X_i$. \square

We will appeal several times throughout this paper to an action of a right-angled Artin group on a set which satisfies the hypotheses of Lemma 3.1. For the sake of completeness and concreteness, we will now exhibit such an action:

Proposition 3.2. *Let Γ be a graph with n vertices. There exists a set X , subsets X_1, \dots, X_n of X and a basepoint $x_0 \in X$ which satisfy the hypotheses of Lemma 3.1.*

Proof. We let X be the set of reduced words in $A(\Gamma)$ and we set x_0 to be the identity. $A(\Gamma)$ acts on this set by left multiplication. We write g_1, \dots, g_n for the vertex generators or right-angled Artin system for $A(\Gamma)$. We let X_i be the set of words of the form w in X which can be written in reduced form as $g_i^s \cdot w'$, where $s \neq 0$ and the length of w is strictly larger than the length of w' . Let $j \neq i$. If (i, j) is not an edge in Γ and $w \in X_i$ then $g_j^k \cdot w$ is clearly contained in X_j for all nonzero k . If (i, j) is an edge in Γ then $g_j^k \cdot w$ is contained in X_i , since we can commute g_j with g_i .

In the latter case, it is possible that $g_j^k w$ is not reduced. However, if any reduction occurs then we must be able to write a reduced expression $w = g_j^t g_i^s w''$ for some nonzero t and some word w'' which is strictly shorter than w' . After we perform all possible reductions to the left of the left-most occurrence of g_i , we can move the occurrences of g_i to the far left. \square

We now give another perspective on the ping-pong lemma for right-angled Artin groups which was pointed out to the author by M. Kapovich and can also be found in [38].

Second proof of Lemma 3.1. Write $w \in A(\Gamma)$ as a reduced word in the vertices of Γ . We call a word **central** if it is a product of vertex generators which commute with each other (in other words, the word is a product of vertices which sit in a complete subgraph of Γ). We say that $w \in A(\Gamma)$ is in **central form** if it is written as a product $w = w_n \cdots w_1$ of central words which is maximal, in the sense that the last letter of w_j does not commute with the last letter of w_{j-1} (where we read from left to right).

Now let x be the basepoint and let $w = w_n \cdots w_1$ be a nontrivial element of $A(\Gamma)$, written in central form. We claim that $w(x)$ is contained in the X_i corresponding to the last letter of w_n . We proceed by induction on n .

If $n = 1$ then conditions 1 and 3 of the assumptions on $A(\Gamma)$ and X gives the claim. Now we consider $w_n(w_{n-1} \cdots w_1(x))$. By induction, $w_{n-1} \cdots w_1(x)$ is contained in the X_i corresponding to the last letter v_i of w_{n-1} . The last letter v_j of w_n does not commute with v_i , so condition 2 implies that $w_n(X_i) \subset X_j$, so that $w(x) \in X_j$. In particular, $w(x) \neq x$. \square

4. ANALOGY WITH FINITE VOLUME HYPERBOLIC k -MANIFOLDS AND RANK ONE LIE GROUPS

In this section we establish parallel results in the theory of hyperbolic k -manifolds and lattices in more general rank one Lie groups. In this section only, we will be using Γ to denote a lattice and Δ to denote a graph since the notation of Γ for a lattice seems to be more standard.

4.1. Real hyperbolic spaces. Let $\Gamma < PSL_2(\mathbb{C})$ be generated by elements $\{g_1, \dots, g_m\}$, no two of which share a fixed point at infinity. If each g_i is a hyperbolic isometry of \mathbb{H}^3 then a standard ping-pong argument proves:

Proposition 4.1. *There exists a positive natural number N such that for all $n \geq N$, we have $\langle g_1^n, \dots, g_m^n \rangle$ is a free group.*

We may assume that every element of a discrete subgroup Γ of $SO(k, 1)$ or $SU(k, 1)$ is either parabolic or hyperbolic, as discreteness allows us to neglect elliptic elements by passing to a finite index subgroup. Furthermore, discreteness guarantees that if $g \in \Gamma$ stabilizes a point x_0 at infinity then either g is contained in a parabolic subgroup Γ_0 which is the stabilizer of x_0 in Γ , or g is hyperbolic and the stabilizer of x_0 is cyclic. A discrete parabolic subgroup of $SO(k, 1)$ is virtually a free abelian group (see any reference on hyperbolic geometry such as the book [4] of Benedetti and Petronio).

Recall that Proposition 1.12 says that Proposition 4.1 holds for general real hyperbolic spaces, even if we allow parabolic isometries as well as hyperbolic isometries.

Proof of Proposition 1.12. Label the fixed points of generators of Γ as

$$A_1, R_1, \dots, A_m, R_m$$

and

$$P_1, \dots, P_{m'},$$

where (A_i, R_i) are the pair of fixed points of a hyperbolic isometry and P_i is the fixed point of a parabolic isometry. After passing to a finite power of the generators of Γ we may assume that for each i , the stabilizer of P_i is a free abelian group. By discreteness, no A_i is equal to any R_j , and the parabolic fixed locus is different from the hyperbolic fixed locus. Choose small neighborhoods V_i of (A_i, R_i) for each i , so that the hyperbolic isometries play ping-pong with each other.

For a fixed i , conjugate P_i to be ∞ . For any norm on \mathbb{R}^k , there is an N such that all other fixed points of generators of Γ have norm less than N . If the stabilizer of ∞ is cyclic and generated by g then there is an M such that $g^{\pm M}(V_i) \subset \{x \mid |x| > N\}$. Thus, g^M plays ping-pong with the hyperbolic isometries. Conjugating back, it follows that for any parabolic isometry g_i stabilizing P_i and any compact subset K of $\partial\mathbb{H}^k \setminus P_i$, there is an $M(K)$ such that for all $m \geq M$, we have $g^m(K) \cap K = \emptyset$.

Suppose that the stabilizer S of ∞ is isomorphic to \mathbb{Z}^s for $s > 1$. We have that S acts on $\mathbb{R}^{k-1} = \partial\mathbb{H}^k \setminus \{\infty\}$ by fixed-point free Euclidean isometries. By discreteness and residual finiteness of finitely generated abelian groups, we see that for any parabolic subgroup S stabilizing P_i and any compact subset $K \subset \partial\mathbb{H}^k$, there is a finite index subgroup S_K of S such that for any $1 \neq g \in S_K$, we have $g(K) \cap K = \emptyset$.

Ping-pong is thus setup by taking the $\{X_i\}$ to be the $\{V_i\}$ for the hyperbolic isometries, and small neighborhoods $\{U_j\}$ of each parabolic fixed point. We choose these neighborhoods to be small enough so that they are pairwise disjoint and do not cover all of the sphere at infinity. The complement of each such neighborhood is a compact subset of $\partial\mathbb{H}^k$. It follows that there is an M such that for all $m \geq M$, the isometry g_i^m maps the compact set $\partial\mathbb{H}^k \setminus X_i$ into X_i . By irredundancy and discreteness, g_i and g_j commute if and only if they generate a parabolic subgroup. If g_i and g_j do not commute then $X_j \subset \partial\mathbb{H}^k \setminus X_i$. Lemma 3.1 gives us the conclusion of the proposition. \square

The requirement that the isometries do not share fixed points at infinity is essential. For instance, the isometries $z \mapsto z + 1$ and $z \mapsto 2z$ of \mathbb{H}^2 have many nontrivial relations, as do their powers. These two isometries generate a solvable group which is not virtually abelian. More generally, the stabilizer of a point at infinity in the group of isometries of \mathbb{H}^2 is a two-step solvable Lie group (the affine group). It can be shown that the only discrete subgroups of this affine group are virtually abelian (cf. S. Katok's book [36]).

4.2. Complex hyperbolic spaces. We now make a short digression into complex hyperbolic geometry so that we can make sense of right-angled Artin systems for discrete subgroups of $SU(k, 1)$ and prove Proposition 1.13. For more details, consult the book [28] of Goldman, for instance. Recall that in analogy to the usual hyperbolic space which we call **real hyperbolic space**, there is a notion of **complex hyperbolic space**. It is often denoted $\mathbb{H}_{\mathbb{C}}^k$. Like real hyperbolic space it can be constructed from a certain level set of the indefinite inner product on $\mathbb{C}^{k,1}$. There is again a notion of a boundary $\partial\mathbb{H}_{\mathbb{C}}^k$ (which is topologically a sphere) to which the isometry group of $\mathbb{H}_{\mathbb{C}}^k$ extends. Isometries are classified analogously to isometries of real hyperbolic space, into hyperbolic, elliptic and parabolic types. If Γ is a finitely generated discrete group of isometries, then Γ virtually contains no elliptic elements.

The stabilizer of a point p in $\partial\mathbb{H}_{\mathbb{C}}^k$ within the isometry group of complex hyperbolic space is identified with a real $(2k - 1)$ -dimensional Heisenberg group H . This group is a 2-step nilpotent real Lie group which admits a real basis

$$x_1, y_1, \dots, x_{k-1}, y_{k-1}, z,$$

satisfying the relations $[x_i, y_i] = z$ for each i , and all other commutators being trivial. The Heisenberg group acts simply transitively on $\partial\mathbb{H}_{\mathbb{C}}^k \setminus p$, so that $\partial\mathbb{H}_{\mathbb{C}}^k$ can be viewed as the one-point

compactification of H . Note that we may assume that the discrete parabolic subgroups of the isometry group of $\mathbb{H}_{\mathbb{C}}^k$ are all torsion-free and at most two-step nilpotent.

Let Γ_0 be a discrete parabolic group of isometries of $\mathbb{H}_{\mathbb{C}}^k$ fixing a point $p \in \partial\mathbb{H}_{\mathbb{C}}^k$. Without loss of generality, Γ_0 is a torsion-free lattice in $H = \partial\mathbb{H}_{\mathbb{C}}^k \setminus p$. Let g_1, \dots, g_m generate Γ_0 . Since Γ_0 is torsion-free, H is isomorphic to the real Mal'cev completion of Γ_0 , so that we may assume that Γ_0 is the group of integer points $H_{\mathbb{Z}}$ of H .

Note that $H_{\mathbb{Z}}$ is residually finite. Discreteness and residual finiteness imply that if K is any compact subset of H and $x \in K$, there is a finite index normal subgroup $N_{x,K} \subset H_{\mathbb{Z}}$ such that if $1 \neq n \in N_{x,K}$, then $n \cdot x \notin K$. In fact, we may choose one such $N_{x,K}$ which works for any $x \in K$. Fixing K , if we choose sufficiently large powers of given elements of $H_{\mathbb{Z}}$, we may assume they lie in $N_{x,K} = N_K$.

Before we give the proof of Proposition 1.13, we state a version of the ping-pong lemma which applies to general free products, and whose proof can be found in [18]:

Proposition 4.2. *Let G be a group acting on a set X , let Γ_1 and Γ_2 be two subgroups of G , and let Γ be the subgroup generated by Γ_1 and Γ_2 . Assume that Γ_1 contains at least three elements and Γ_2 contains at least two elements. Assume furthermore that there exist two non-empty subsets X_1 and X_2 of X with X_2 not included in X_1 such that $\gamma_1(X_2) \subset X_1$ for every nonidentity $\gamma_1 \in \Gamma_1$ and $\gamma_2(X_1) \subset X_2$ for every nonidentity $\gamma_2 \in \Gamma_2$. Then $\Gamma \cong \Gamma_1 * \Gamma_2$.*

We are now ready to prove Proposition 1.13.

Proof of Proposition 1.13. Let $\Gamma < SU(k, 1)$ be discrete. Without loss of generality, Γ is torsion-free. Let Γ be generated by $\{g_1, \dots, g_m\}$, which we identify with an irredundant collection of isometries of $\mathbb{H}_{\mathbb{C}}^k$ via the projection

$$SU(k, 1) \rightarrow \text{Isom}(\mathbb{H}_{\mathbb{C}}^k).$$

We group these isometries according to whether they are parabolic or hyperbolic, and then according to which points they stabilize in $\partial\mathbb{H}_{\mathbb{C}}^k$. We claim that for n sufficiently large, we have

$$\langle g_1^n, \dots, g_m^n \rangle \cong F_{\ell} * N_1 * \dots * N_s,$$

where ℓ is the number of distinct hyperbolic geodesics stabilized by elements of $\{g_1, \dots, g_m\}$, and each N_i is a subgroup of the stabilizer of some parabolic fixed point of certain elements in the collection $\{g_1, \dots, g_m\}$.

Fixed points of hyperbolic isometries come in pairs, one element of which is attracting and the other of which is repelling. It follows that we can find neighborhoods of these fixed points and powers of the hyperbolic generators of Γ so that they play ping-pong with each other.

Let N be a discrete parabolic subgroup of $\text{Isom}(\mathbb{H}_{\mathbb{C}}^k)$ and let $K \subset \partial\mathbb{H}_{\mathbb{C}}^k$ be a compact set which does not contain the fixed point p of N . Let U be a neighborhood p which is disjoint from K . By the residual finiteness of N , we can find a finite index subgroup $N' < N$ such that every nonidentity element of N' sends K into U . If $g_i \in N$ is a generator of Γ then there is an n such that $g_i^n \in N'$.

Choose disjoint neighborhoods of all the fixed points of elements in the collection $\{g_1, \dots, g_m\}$ and an n sufficiently large such that if U and V are neighborhoods of the distinct fixed point(s) of g and h respectively, $g^{\pm n}(V) \subset U$ and $h^{\pm n}(U) \subset V$. This is possible for hyperbolic and parabolic isometries by the previous two paragraphs. But then the hypotheses of Lemma 4.2 are satisfied, giving the desired conclusion. \square

4.3. Commensurability and embeddings. We now return to the world of real hyperbolic manifolds. Let Γ be the fundamental group of a finite volume hyperbolic 3-manifold. If \mathbb{H}^3/Γ is closed then there are no noncyclic abelian subgroups of Γ . Thus if Γ is virtually isomorphic to a right-angled Artin group then Γ is virtually free. This violates Mostow rigidity (see [4]), since free groups have infinite outer automorphism groups.

We will first determine the right-angled Artin groups which do embed as discrete groups of isometries of real hyperbolic spaces. In [34], Januszkiewicz and Świątkowski study Coxeter groups of high virtual cohomological dimension. They prove certain dimension restrictions on Coxeter groups which are also virtual Poincaré duality groups in both the right-angled and non-right-angled cases. Along similar lines, we can characterize the right-angled Artin groups which admit a discrete, faithful representation into $SO(k, 1)$, which we call **hyperbolic right-angled Artin groups**.

Proposition 4.3. *The hyperbolic right-angled Artin groups are precisely the finite free products of free abelian groups.*

Proof. By choosing k sufficiently large, it is clear that a finite free product of free abelian groups can be embedded into $SO(k, 1)$. Conversely, suppose $A(\Delta)$ is a hyperbolic right-angled Artin group. Let A_N be the isomorphism type of the group given by replacing each element of the right-angled Artin system by its N^{th} power. It is evident that there is an isomorphism $A(\Delta) \cong A_N$, and this can be verified by a ping-pong calculation. If the elements of the right-angled Artin system for $A(\Delta)$ do not already play ping-pong at infinity as in the proof of Proposition 1.12, we may replace them by a sufficiently large power so that they do. It follows that $A(\Delta)$ has to be a finite free product of free abelian groups. \square

I. Agol and D. Wise have recently announced a proof that hyperbolic 3-manifolds are virtually special and hence virtually inject into right-angled Artin groups. However:

Proposition 4.4. *Let $\Gamma < \text{Isom}(\mathbb{H}^3)$ be a lattice. Then Γ is not commensurable with a right-angled Artin group.*

Proof. After passing to a finite index subgroup, we may assume there is no torsion in Γ . Consider $1 \neq g \in \Gamma$. Since the elements of Γ which centralize g must have the same fixed points on $\widehat{\mathbb{C}}$ as g , we have that the centralizer $C_\Gamma(g) \cong \mathbb{Z}$ when g is hyperbolic or $C_\Gamma(g) \cong \mathbb{Z}^2$ when g is parabolic. If Γ is commensurable to a right-angled Artin group $A(\Delta)$ then $A(\Delta)$ cannot have abelian subgroups of rank greater than two. In particular Δ cannot contain a complete graph on three vertices.

If h, h' both centralize with g then both h and h' have the same fixed points as g and thus commute with each other. It follows that $C_\Gamma(g)$ is abelian. In particular, Δ cannot have a non-backtracking path of length more than two. It follows that $A(\Delta)$ splits as a free product of finitely many copies of \mathbb{Z}^2 . By the work of Droms in [19] (or by the Kurosh Subgroup Theorem), every finitely generated subgroup of $A(\Delta)$ is a right-angled Artin group, so that Γ is virtually a right-angled Artin group. The previous considerations show that Γ must split as a nontrivial free product of cyclic groups and copies of \mathbb{Z}^2 . Any such free product has an infinite outer automorphism group, which again violates Mostow Rigidity. \square

We would like to apply Mostow Rigidity to establish Proposition 4.4 for a more general class of hyperbolic manifolds. Before giving the proof of Theorem 1.14, we first make a few remarks. Let X be a right-angled Artin system for a right-angled Artin group. Let α be a **Whitehead automorphism** of a right-angled Artin group, which is to say one which restricts to $X \cup X^{-1}$ as one of the following:

- (1) A permutation of $X \cup X^{-1}$.
- (2) There is an $a \in X$ such that for each $x \in X$,

$$\alpha(x) \in \{x, ax, xa^{-1}, axa^{-1}\}.$$

It is known by the work of M. Laurence in [39] that Whitehead automorphisms generate the whole automorphism group of the right-angled Artin group (see also the work of M. Day in [17]).

Let Γ be a lattice in $SO(k, 1)$ and let $A(\Delta)$ be a candidate right-angled Artin group which is commensurable with Γ . We may assume that there is a finite index subgroup G of $A(\Delta)$ such

that $\Gamma \cong G$, since a finite index subgroup of a lattice is still a lattice. If $n \geq 3$ then Mostow Rigidity implies that $\text{Out}(\Gamma)$ is finite. Whenever there exists an infinite order, non-inner Whitehead automorphism, it can be used to show that $\text{Out}(G)$ is in fact infinite, which gives us a contradiction.

One example of a Whitehead automorphism of a right-angled Artin group $A(\Delta)$ is a **partial conjugation**. Let x be a vertex of Δ . One conjugates the generators of $A(\Delta)$ corresponding to the vertices in a (union of) connected component of the complement of the star of x by the generator corresponding to x . Whenever x separates Δ , this gives a non-inner automorphism of $A(\Delta)$ with infinite order in $\text{Out}(A(\Delta))$.

Another example of a Whitehead automorphism is a **dominated transvection**. We say a vertex w dominates a vertex v whenever the link of v is contained in the star of w . When w dominates v , the map $v \mapsto vw$ is an infinite order non-inner automorphism of the right-angled Artin group, and is called a dominated transvection. It is clear that this transvection acts nontrivially on the homology of the right-angled Artin group. It follows that a dominated transvection has infinite order in $\text{Out}(A(\Delta))$.

Usually, the outer automorphism group of a right-angled Artin group admits no non-inner partial conjugacies and no dominated transvections. As a consequence, one can show that usually the outer automorphism group of a right-angled Artin group is finite, as is shown in the recent paper of R. Charney and M. Farber [12]. However, the class of hyperbolic right-angled Artin groups is very restricted, and we will see that hyperbolic right-angled Artin groups are far from generic.

Proof of Theorem 1.14. Suppose that a lattice Γ in $SO(k, 1)$ is commensurable with a right-angled Artin group $A(\Delta)$. Then $A(\Delta)$ must virtually embed as a discrete subgroup of $SO(k, 1)$, in which case it is a free product of free abelian groups. By the Kurosh Subgroup Theorem, it follows that Γ splits as a free product of free abelian groups. Any such group admits an infinite group of outer automorphisms (coming from dominated transvections, for instance), a violation of Mostow rigidity. \square

Note that Theorem 1.14 is absolutely false for $k = 2$. Indeed, $PSL_2(\mathbb{Z})$ is virtually a right-angled Artin group and is also a lattice in $PSL_2(\mathbb{R})$. The reason for this dramatic failure is the fact that maximal parabolic subgroups and stabilizers of hyperbolic geodesics in $PSL_2(\mathbb{Z})$ are indistinguishable from the point of view of their abstract isomorphism type. Since Mostow rigidity works in dimension three and higher, it is evident that the proof above does not work for $k = 2$. We can give an alternative proof of Theorem 1.14 which does not use Mostow rigidity so heavily.

Second proof of Theorem 1.14. Note that we may assume that \mathbb{H}^k/Γ is not closed. If Γ is a lattice in $SO(k, 1)$, we may assume that the maximal rank of a free abelian group in Γ is $k - 1$. Furthermore, any abelian subgroup of Γ of rank $k - 1$ is conjugate into a peripheral subgroup (this is where the argument fails for $k = 2$). It follows that any candidate right-angled Artin group $A(\Delta)$ has cohomological dimension at most $k - 1$, which is to say that its Salvetti complex is built out of tori of dimension no larger than $k - 1$. We are assuming that Γ is not cocompact, so we may double it over the cusps. The resulting group $D(\Gamma)$ is the fundamental group of a closed, orientable k -manifold. Therefore, $H^k(D(\Gamma), \mathbb{Z}) \cong \mathbb{Z}$. On the other hand, we can construct a complex with fundamental group $D(\Gamma)$ out of two copies of a finite cover of the Salvetti complex of $A(\Delta)$. The resulting complex has cohomological dimension at most $k - 1$, a contradiction. \square

Given the discussion in this section, Theorem 1.1 should be thought of as a rank one phenomenon of the mapping class group. It is well-known that mapping class groups share many properties with lattices in rank one semisimple Lie groups as well as properties of lattices in higher rank semisimple Lie groups (see the paper [22] of Farb, Lubotzky and Minsky, for instance).

5. THE BAD BEHAVIOR OF RIGHT-ANGLED ARTIN GROUPS

The purpose of this section is to give examples of badly-behaved right-angled Artin groups, to show that Theorem 1.1 is as general as possible, and to record some other useful facts about right-angled Artin groups. A much more detailed discussion of some of the issues arising here can be found in [37].

In our analysis of mapping classes and their interactions, it will be helpful to write mapping classes as compositions of commuting mapping classes, so that we only need to consider the “simplest” types of mapping classes. In this section, we will consider certain natural subgroups of right-angled Artin groups generated by clique subgroups. We will be able to give sufficient (though not necessary: see [37]) conditions for such a subgroup to again be a right-angled Artin group (see Lemma 5.2 below). After this section, we will restrict our considerations to the case where each mapping class is a Dehn twist about a simple closed curve or a pseudo-Anosov homeomorphism when restricted to a connected subsurface.

If Γ is a graph, we call a nonempty complete subgraph of Γ a **clique** of Γ . The subgroup generated by those vertices is called a **clique subgroup**. A **join** of two graphs is a union of two graphs Γ_1 and Γ_2 , together with edges connecting every vertex of Γ_1 with every vertex of Γ_2 . A **join subgroup** of $A(\Gamma)$ is a subgroup generated by the vertices of a subgraph of Γ which decomposes as a nontrivial join, considered up to conjugacy.

In a right-angled Artin group, commutation of elements is more subtle than containment in a conjugate of a clique subgroup. Behrstock and Charney prove the following result, which is very useful in understanding commutation in right-angled Artin groups:

Proposition 5.1 ([3]). *Let $g \in A(\Gamma)$ be cyclically reduced. The following are equivalent:*

- (1) $g \in A(\Gamma)$ is contained in a join subgroup.
- (2) The centralizer $C(g)$ is contained in a join subgroup.
- (3) The centralizer of g is not cyclic.

Let Γ be a finite graph and let $A(\Gamma)$ be the corresponding right-angled Artin group. Identify the vertices of Γ with a right-angled Artin system for $A(\Gamma)$ and hence a set of generators for $A(\Gamma)$. Let C_1, \dots, C_n be the cliques of Γ , and let $\{z_{i,j}\}$, $1 \leq i \leq n$ and $1 \leq j \leq m(i)$ be (possibly empty) collections of (formal positive powers of) vertices (or their formal inverses) in $\{C_i\}$. Let

$$\{Z_i = \prod_j z_{i,j}\},$$

and form a graph Λ from the collection $\{Z_i\}$ by taking one vertex for each generator and connecting two vertices if the two elements commute in $A(\Gamma)$. We will say that a collection $\{Z_i\}$ has **property PP** (which stands for ping-pong) if Λ occurs as an induced subgraph of Γ in a way such that the vertex v_i of Λ corresponding to Z_i lies in the support of Z_i under the inclusion $\Lambda \rightarrow \Gamma$. Here, a subgraph $\Lambda \subset \Gamma$ is **induced** if two vertices in Λ are adjacent if and only if they are adjacent in Γ .

Lemma 5.2. *Let $\{Z_i\}$ have property PP. Then*

$$\langle \{Z_i\} \rangle \cong A(\Lambda).$$

Proof. There is clearly a surjective map $A(\Lambda) \rightarrow \langle \{Z_i\} \rangle$, which we claim is also injective. Make $A(\Gamma)$ act faithfully on a set X , and for each vertex z of Γ , let $X_z \subset X$ be the subset which witnesses the ping-pong lemma for the action. Since Λ is an induced subgraph of Γ , let z_i be the image of the vertex v_i under the inclusion $\Lambda \rightarrow \Gamma$. Write each Z_i and its inverse in such a way so that all of the vertex letters z_i corresponding to v_i are as far to the right as possible. Let x_0 be the basepoint of the ping-pong action. Clearly for all nonzero powers of Z_i , the image of x_0 is in X_{v_i} .

If v_i and v_j are connected in Λ then the vertices in the supports of Z_i and Z_j form one large clique (this fact is not entirely obvious but can be deduced from Proposition 5.1. See [37] for more details). Therefore, each nonzero power of Z_j preserves X_{v_i} .

If v_i and v_j are not connected in Λ then the vertices z_i and z_j are not connected in Γ . It follows that each nonzero power of Z_j sends X_{v_i} into X_{v_j} . This verifies the ping-pong lemma and proves the result. \square

We now illustrate some examples of subgroups of right-angled Artin groups which one might expect to be right-angled Artin groups themselves but which are not, or which are right-angled Artin groups of an unexpected isomorphism type.

Let A, B, C generate a copy of \mathbb{Z}^3 and let B, C, X, Y generate a copy of \mathbb{Z}^4 . Let $A(\Gamma)$ be the corresponding right-angled Artin group generated by A, B, C, X, Y , which is to say there are no relations between A and $\langle X, Y \rangle$. Consider the subgroup G generated by AB, AC, X, Y . One checks easily that AB and AC both do not commute with X and Y . The naïve guess as to the isomorphism type of G would be $\mathbb{Z}^2 * \mathbb{Z}^2$. This is not right, however. Indeed, $(AC)^{-1}AB = BC^{-1}$, which commutes with both X and Y . It follows that X and Y do not generate a maximal abelian subgroup, but rather X, Y, BC^{-1} do. It follows that if G is a right-angled Artin group, it must be the one whose underlying graph Γ' is a triangle with a single edge emanating out of one of the vertices. The vertices should be labeled X, Y, BC^{-1} on the triangle, and BC^{-1} should be connected to AC .

Proposition 5.3. *There is an isomorphism $G \cong A(\Gamma')$.*

Proof. It is crucial that we already know that $A(\Gamma)$ is a right-angled Artin group. It is easy to check that the vertex C dominates both A and B . It follows that the dominated transvections $A \mapsto AC^{\pm 1}$ and $B \mapsto BC^{\pm 1}$ give automorphisms of $A(\Gamma)$. It follows that the group generated by $\{X, Y, BC^{-1}, AC\}$ is isomorphic to the subgroup generated by $\{X, Y, B, A\}$, which is precisely $A(\Gamma')$. \square

We will now give an example illustrating the fact that given an arbitrary collection of mapping classes, there might be no powers which generate a right-angled Artin group. Let c_1, c_2, c_3, c_4 be simple closed curves on a surface Σ , with

$$i(c_1, c_3), i(c_2, c_4) \neq 0,$$

c_2 disjoint from both c_1 and c_3 , c_1 disjoint from c_4 , and c_3 disjoint from c_4 . Let T_1, T_2, T_3 and T_4 be Dehn twists about these curves. The proof of Theorem 1.1 will show that there are powers w, x, y and z of T_1, T_2, T_3 and T_4 respectively which generate a right-angled Artin group, and the corresponding graph will be a square, so that $\langle w, x, y, z \rangle \cong F_2 \times F_2$. When it makes sense, we will write the composition of these mapping classes additively.

Proposition 5.4. *The eight mapping classes $w - x, w + x, x - y, x + y, y - z, y + z, w - z$ and $w + z$ have the property that no nonzero powers generate a right-angled Artin group, in the sense that they are not a right-angled Artin system for any right-angled Artin group.*

Proof. Let A_1 be the group generated by $w - x$ and $w + x$, let A_2 be the group generated by $x - y$ and $x + y$, let A_3 be the group generated by $y - z$ and $y + z$, and let A_4 be the group generated by $w - z$ and $w + z$. These are all isomorphic to \mathbb{Z}^2 . If powers of these mapping classes generated a right-angled Artin group, it would be one on a finite list of possible isomorphism types. We can immediately rule out \mathbb{Z}^4 , any group which contains a copy of \mathbb{Z}^3 , and any right-angled Artin group on a graph with fewer than four edges. The only possible underlying graph is the square.

Consider the intersection of A_1 and A_2 . It is easy to see that the intersection is generated by $2x$. Similarly, the intersection of A_2 and A_3 is generated by $2y$. Note that $2x$ and $2y$ together do not generate A_2 , since

$$[\langle x, y \rangle : A_2] = 2$$

and

$$[\langle x, y \rangle : \langle 2x, 2y \rangle] = 4.$$

Suppose we replace $x - y$ and $x + y$ by $n(x - y)$ and $m(x + y)$, and we find the smallest positive integers a and b such that $an = bm$. Then $\langle n(x - y), m(x + y) \rangle$ contains $(an + bm)x$ and $(an - bm)y$.

Note that $\langle x - y, x + y \rangle$ only contains even multiples of x and y . Therefore, if we want $x + y$ and $x - y$ to be contained in the group generated by the vertex generators, we may assume that both m and n are even but not divisible by four. But then $(an + bm)$ and $(an - bm)$ are both divisible by four, since $\langle x - y, x + y \rangle$ only contains even multiples of x and y . Alternatively, since $n = 2n'$ and $m = 2m'$ with both n' and m' odd, we can check case by case that if $an' = bm'$ then $an' + bm'$ and $an' - bm'$ are both even. Thus, $A_{n,m} = \langle n(x - y), m(x + y) \rangle$ always properly contains

$$\langle A_{n,m} \cap \langle x \rangle, A_{n,m} \cap \langle y \rangle \rangle.$$

We can repeat this argument for A_1 , A_3 and A_4 . Let n_1, \dots, n_4 and m_1, \dots, m_4 be the multiples by which we replace $w - x, x - y, y - z, w - z$ and $w + x, x + y, y + z, w + z$ respectively. We have that $\langle n_1(w - x), m_1(w + x) \rangle$ properly contains the subgroup generated by its intersection with the groups $\langle w \rangle$ and $\langle x \rangle$. It is possible that

$$\langle n_2(x - y), m_2(x + y) \rangle$$

and

$$\langle n_4(w - z), m_4(w + z) \rangle$$

correct this deficiency by containing smaller multiples of w and x , but in that case at least one of

$$\langle n_2(x - y), m_2(x + y) \rangle$$

or

$$\langle n_4(w - z), m_4(w + z) \rangle$$

properly contains the group generated by its intersection with the corresponding vertex groups. It may be possible that

$$\langle n_3(y - z), m_3(y + z) \rangle$$

corrects the deficiency, but then

$$\langle n_3(y - z), m_3(y + z) \rangle$$

properly contains the group generated by its intersection with $\langle y \rangle$ and $\langle z \rangle$. \square

The last part of this section serves as a motivation for the definition of property PP and an illustration of some of the bad behavior of right-angled Artin groups.

The author is indebted to M. Casals for the following example. Consider the group $F_2 \times F_2$ generated by w, x, y, z as above. The elements w, x, yz are not a right-angled Artin system for any right-angled Artin group, and neither are any of their powers. Denoting yz by g , it is easy to check that w commutes with the conjugate of x by g , and the same relation holds for all powers of w, x , and g .

Proposition 5.5. *Consider the group G_N generated by w^N, x^N and g^N as above, for $N \neq 0$. Then G_N is not even abstractly isomorphic to a right-angled Artin group.*

Proof. We will prove the proposition for $G = G_1$. The proof in general is exactly the same. We claim that G is not isomorphic to any right-angled Artin group, for which it suffices to check that G is not $\mathbb{Z}^2 * \mathbb{Z}$ or $F_2 \times \mathbb{Z}$ since G is obviously neither free nor abelian. For the second case, it is trivial to understand the centralizer of any element and to show that G has no center.

Suppose that G splits as a free product of \mathbb{Z}^2 and \mathbb{Z} . The classical Kurosh Subgroup Theorem implies that any copy of \mathbb{Z}^2 inside of G is conjugate to a subgroup of the \mathbb{Z}^2 factor of the splitting. Alternatively, one can see this by constructing the Cayley graph of $\mathbb{Z}^2 * \mathbb{Z}$ with the standard generating set and noticing that any copy of the \mathbb{Z}^2 Cayley graph must come from a conjugate of the copy based at the identity. In G , $\langle w, x \rangle$ and $\langle w, x^g \rangle$ are both copies of \mathbb{Z}^2 , and it is easy to check that these two copies of \mathbb{Z}^2 are not conjugate in G . \square

6. RIGHT-ANGLED ARTIN GROUPS AS SUBGROUPS OF MAPPING CLASS GROUPS

The goal of this section is to prove Theorem 1.1. To start out, we recall some background information and make some preliminary simplifying remarks. A reader familiar with the basics of mapping class groups and the action of $\text{Mod}_{g,p}$ on $\mathbb{P}\mathcal{ML}(\Sigma)$ may wish to skip to subsection 6.6.

6.1. Reducible mapping classes. In the course of proving our main theorem, we may certainly replace $\text{Mod}_{g,p}$ with any finite index subgroup. In particular, we may assume that all the mapping classes we consider are **pure**. Recall the description of pure mapping classes:

Lemma 6.1. *Let f be a pure mapping class. Then there exists a system of disjoint, essential, pairwise non-isotopic curves*

$$\mathcal{C} = \bigcup \gamma_i$$

such that:

- (1) \mathcal{C} is preserved component-wise by f .
- (2) $\Sigma \setminus \mathcal{C}$ consists of a union of simpler surfaces, each of which is preserved component-wise by f .
- (3) The restriction of f to any component of $\Sigma \setminus \mathcal{C}$ is either pseudo-Anosov or trivial.

The description of pure mapping classes in Lemma 6.1 allows us to characterize (virtual) centralizers in $\text{Mod}_{g,p}$. Virtually solvable subgroups of $\text{Mod}_{g,p}$ were classified in [6], where it was shown that they are virtually abelian. We will be requiring an in depth understanding of virtually abelian subgroups of mapping class groups here. To that end, we need to define a **canonical reduction system** for a mapping class. For each pure mapping class f , we consider the collection \mathcal{C} of all simple closed curves whose isotopy class is preserved by f . Let $\{\mathcal{C}_i\}_{i \in I}$ be the set of maximal (though possibly empty) collections of disjoint simple closed curves which witness the fact that f is a pure mapping class. The canonical reduction system \mathcal{C}_R is defined by

$$\mathcal{C}_R = \bigcap_{i \in I} \mathcal{C}_i.$$

We make a few elementary observations. Let Σ_j be a component of $\Sigma \setminus \mathcal{C}_i$. If the restriction f_j of f to Σ_j is pseudo-Anosov relative to its boundary components, then Σ_j cannot be reduced any further. Notice also that if \mathcal{C}_k is any other reduction system for f , then the boundary components of Σ_j are contained in \mathcal{C}_k . Indeed, otherwise there would be a component of $\Sigma \setminus \mathcal{C}_k$ which properly contains Σ_j for which the inclusion is not a homotopy equivalence, so that \mathcal{C}_k is not a reduction system. Furthermore, let $c \subset \mathcal{C}_i$ be a simple closed curve. We have that c contributes two boundary components to one or two components of $\Sigma \setminus \mathcal{C}_i$. Suppose that the restriction of f to these components is the identity and that the restriction to their union together with c is not a nontrivial power of a Dehn twist about c . Then $\mathcal{C}_i \setminus \{c\}$ is also a reduction system for f , so that $c \notin \mathcal{C}_R$.

The following can be found in Ivanov's book [33], for instance. We recall the proof here for the reader's convenience.

Lemma 6.2. *Let f and g be two nontrivial pure mapping classes and let \mathcal{C}_1 and \mathcal{C}_2 be (possibly empty) canonical reduction systems for f and g with restrictions $\{f_i\}$ and $\{g_i\}$ to the interiors of the components of $\Sigma \setminus \mathcal{C}_i$. Then f and g virtually commute if and only if one of the following conditions holds:*

- (1) f and g are both pseudo-Anosov and share a common nonzero power. This happens if and only if f and g stabilize the same geodesic in Teichmüller space.
- (2) $\mathcal{C}_1 = \mathcal{C}_2$ and the restrictions f_i and g_i for each i are virtually commuting pseudo-Anosov mapping classes, or one of the two restrictions is trivial. Note that this case encompasses the first case when the reduction systems are empty.

- (3) \mathcal{C}_1 is properly contained in \mathcal{C}_2 and if Σ_i is a component of $\Sigma \setminus \mathcal{C}_1$ which contains at least one component of \mathcal{C}_2 , then the restriction of f to Σ_i is trivial.
- (4) $\mathcal{C}_1 \cup \mathcal{C}_2$ form part of a pants decomposition of Σ , and if $c \in \mathcal{C}_1$ is contained in a component $\Sigma_i \subset \Sigma \setminus \mathcal{C}_2$, the restriction g_i of g to Σ_i preserves c , and conversely after switching the roles of f and g .

Proof. Clearly the “if” direction holds.

If both f and g are pseudo-Anosov and do not stabilize the same Teichmüller geodesic, then by replacing f and g by positive powers we obtain a free group (a more general statement is proved in Lemma 6.3).

Suppose $\mathcal{C}_1 \subset \mathcal{C}_2$ and $\Sigma_i \subset \Sigma \setminus \mathcal{C}_1$ contains at least one component of \mathcal{C}_2 , and suppose that the restriction of f to Σ_i is pseudo-Anosov. There is a collection of simple closed curves $\mathcal{C}' \subset \Sigma_i$ which is preserved by g but not by f . If $c \in \mathcal{C}'$ then some N we have that the geometric intersection number of $f^n(c)$ and c is nonzero for $n > N$, since $f^n(c)$ converges to the stable lamination of f . In particular, a Dehn twist about c will not commute with f^n . More generally, applying f to \mathcal{C}' does not preserve it, so conjugating g by f will not give back g .

Finally, suppose that there are no inclusion relations between \mathcal{C}_1 and \mathcal{C}_2 . Then either there exist a pair $c_i \in \mathcal{C}_i$ with positive geometric intersection or not. If not, each $c \in \mathcal{C}_1$ sits non-peripherally in $\Sigma \setminus \mathcal{C}_2$, and similarly after switching the indices. At least one of the restrictions in $\{f_i\}$ or $\{g_i\}$ is nontrivial, or f and g are powers of Dehn twists about curves in \mathcal{C}_1 and \mathcal{C}_2 . We may assume that $c \in \mathcal{C}_1$ sits in $\Sigma_i \subset \Sigma \setminus \mathcal{C}_2$ and that the restriction g_i is pseudo-Anosov. It follows that f and g cannot commute.

Suppose that there is an intersecting pair of simple closed curves in the two reduction systems. We have that in $\Sigma \setminus \mathcal{C}_1$, c_1 contributes two boundary components, and the restriction of f to the interior of the components of $\Sigma \setminus \mathcal{C}_1$ which contains the components of c_1 is either pseudo-Anosov or trivial. If it is trivial, then by our conventions on canonical reducing systems we must have that f does a nonzero power of a Dehn twist about c_1 . Since c_2 has positive intersection number with c_1 , we have that g conjugated by f cannot preserve c_2 . Therefore we may assume that the restriction is pseudo-Anosov on at least one of the components. Call such a component X . We have that c_2 intersects X is a union of arcs \mathcal{A} . Since X itself is a surface of hyperbolic type and the restriction of f to X is pseudo-Anosov, it follows that $f(\mathcal{A}) \neq \mathcal{A}$, even if modify f by arbitrary Dehn twists along curves in \mathcal{C}_1 . It follows that $f(c_2) \neq c_2$, so that f and g do not commute. \square

We are in a position to establish the following well-known special case of Theorem 1.1:

Lemma 6.3. *Let $\{f_1, \dots, f_k\} \in \text{Mod}_{g,p}$ be pseudo-Anosov mapping classes. Then there exists an N such that for all $n \geq N$, the subgroup of $\text{Mod}_{g,p}$ generated by $\{f_1^n, \dots, f_k^n\}$ is a free group of rank $r \leq k$. Equality holds if and only if no pair $\{f_i, f_j\}$ generates a virtually cyclic group for $i \neq j$.*

Proof. This is a consequence of the ping-pong argument on $\mathbb{P}\mathcal{ML}(\Sigma)$, and is fleshed out in more detail in [33]. If $f_i^n \neq f_j^m$ for all $n, m \neq 0$, then they stabilize distinct pairs of measured laminations in $\mathbb{P}\mathcal{ML}(\Sigma)$. The induced dynamics of each pseudo-Anosov map on $\mathbb{P}\mathcal{ML}(\Sigma)$ is to attract to the stable lamination and repel from the unstable lamination. Outside of a small neighborhood of both the stable and unstable laminations, the attraction/repulsion is uniform, by compactness.

Let V_i^u and V_i^s be small neighborhoods about the unstable and stable laminations of f_i in $\mathbb{P}\mathcal{ML}(\Sigma)$. For a sufficiently large n_i , f_i will map all of $\mathbb{P}\mathcal{ML}(\Sigma) \setminus V_i^u$ into V_i^s . The ping-pong argument now applies to show that $\{f_i^{n_i}\}$ generate a free group. If $f_i^n = f_j^m$ for some n and m , we replace f_j by f_j^m and delete f_i from our list of mapping classes. We can take $N = \max n_i$. \square

At this point it is important to remark about a subtle distinction between pseudo-Anosov mapping classes and hyperbolic isometries of \mathbb{H}^k . On the one hand, by the results of the exposés of Fathi, Laudenbach and Poénaru in [24], a single pseudo-Anosov lamination determines the pseudo-Anosov homeomorphism, up to a power. On the other hand, a fixed point in the sphere at infinity

S_∞ does not uniquely determine the axis of a hyperbolic isometry of \mathbb{H}^k , and we have noted before that ping-pong fails for pairs of hyperbolic isometries which share exactly one fixed point at infinity. We never encounter this problem with pseudo-Anosov homeomorphisms.

Let f be a pure mapping class with canonical reduction system \mathcal{C} . We say that f has a **pseudo-Anosov component** if the restriction of f to a component of $\Sigma \setminus \mathcal{C}$ is pseudo-Anosov. Let f and g be pure mapping classes with pseudo-Anosov components and canonical reduction systems \mathcal{C}_1 and \mathcal{C}_2 . We say that the two pseudo-Anosov components have **essential overlap** if whenever \mathcal{C}_1 and \mathcal{C}_2 are arranged to have minimal self-intersection (for instance representing the elements of \mathcal{C}_1 and \mathcal{C}_2 by geodesics in some hyperbolic metric), there are pseudo-Anosov components of f and g which intersect nontrivially. We say that the essential overlap is **distinct** if when the two pseudo-Anosov components coincide, the two restrictions do not share common power.

To prove Theorem 1.1, it will be necessary to study the dynamics of mapping classes acting on projective measured laminations and the interactions between the limiting laminations of the mapping classes in question.

6.2. Convergence in $\mathbb{P}\mathcal{ML}(\Sigma)$. The following facts are both well-known and useful for our purposes, but we will not be providing proofs:

Proposition 6.4. *Let $\epsilon > 0$, and let T be a Dehn twist about a curve c , which we view as an element of $\mathbb{P}\mathcal{ML}(\Sigma)$. Let $K \subset \mathbb{P}\mathcal{ML}(\Sigma) \setminus c$ be a compact set contained in the open set consisting of laminations \mathcal{L} with $i(c, \mathcal{L}) > \epsilon$. Then for any neighborhood U of c which is disjoint from K and any sufficiently large N (which depends on K and U), we have $T^N(K) \subset U$.*

Proposition 6.5. *Let $\epsilon > 0$, and let ψ be a pseudo-Anosov supported on a subsurface of Σ with stable and unstable laminations \mathcal{L}^+ and \mathcal{L}^- respectively. Let $K \subset \mathbb{P}\mathcal{ML}(\Sigma) \setminus \mathcal{L}^-$ be a compact set contained in the open set consisting of laminations \mathcal{L}' with $i(\mathcal{L}^\pm, \mathcal{L}') > \epsilon$. Then for any neighborhood U of \mathcal{L}^+ which is disjoint from K and any sufficiently large $N > 0$ (which depends on K and U), we have $\psi^N(K) \subset U$.*

6.3. Intersections and freeness. In this subsection, we will prove a weak version of Theorem 1.1, which illustrates the fact that intersections of supports of mapping classes give rise to free subgroups of mapping class groups. Let f_i be a mapping class which is either a pseudo-Anosov homeomorphism supported on a connected subsurface Σ_i of Σ or a Dehn twist about a simple closed curve c_i . In the first case, f_i has a pair \mathcal{L}_i^\pm of invariant laminations on Σ_i and in the second case f_i attracts any lamination which intersects c_i to c_i . In this sense, the laminations \mathcal{L}_i^\pm or c_i can be thought of as the **limiting laminations** of f_i .

The following theorem is the main result of this subsection, and is well-known by the work of McCarthy and Ivanov:

Theorem 6.6. *Let $\{f_1, \dots, f_k, t_1, \dots, t_\ell\} \subset \text{Mod}_{g,p}$, with each f_i pseudo-Anosov on a subsurface Σ_i with limiting laminations \mathcal{L}_i^\pm and each t_i a Dehn twist with limiting lamination c_i . Suppose that for each distinct pair of classes in $\{f_1, \dots, f_k, t_1, \dots, t_\ell\} \subset \text{Mod}_{g,p}$, the limiting laminations have essential intersection. Then there exists an $N > 0$ such that for all $n \geq N$, we have $\{f_1^n, \dots, f_k^n, t_1^n, \dots, t_\ell^n\}$ generates a free group.*

Proof. For each limiting lamination \mathcal{L} of a mapping class on our list, write $Z_\mathcal{L}$ for the closed subset of $\mathbb{P}\mathcal{ML}(\Sigma)$ which do not intersect \mathcal{L} . If \mathcal{L} is a limiting lamination of any of the mapping classes on the list, then all the other limiting laminations of the other mapping classes are contained in a compact subset $K \subset \mathbb{P}\mathcal{ML}(\Sigma) \setminus Z_\mathcal{L}$. The result follows immediately from the discussion in Subsection 6.2 and the ping-pong lemma. \square

We remark briefly that Theorem 6.6 is false if we replace $\{f_1, \dots, f_k, t_1, \dots, t_\ell\} \subset \text{Mod}_{g,p}$ by an arbitrary finite collection of pure mapping classes. Consider five Dehn twists T_1, \dots, T_5 such that $[T_1, T_2] = 1$, $[T_3, T_4] = 1$, $[T_1, T_3] = 1$, $[T_1, T_4] = 1$, $[T_2, T_4] = 1$, $[T_2, T_5] = 1$ and $[T_3, T_5] = 1$, and

such that all other pairs of twists do not commute. Write $\psi_1^m, \psi_2^m, \psi_3^m$ for the mapping classes $(T_1T_2)^m, (T_3T_4)^m$ and T_5^m respectively. For each m , we have that

$$(\psi_3^m)^{\psi_1^m} = (\psi_3^m)^{(\psi_1^m)^{\psi_2^m}}.$$

Thus, $\{\psi_1^m, \psi_2^m, \psi_3^m\}$ cannot generate a free group of rank three, even though their supports mutually intersect and though the abelianization of the group they generate is isomorphic to \mathbb{Z}^3 .

6.4. Competing convergence. In this subsection, we discuss the primary technical difficulty which arises when we encounter two disjoint limiting laminations. Suppose \mathcal{L} is a lamination which intersects two disjoint curves c_1, c_2 with corresponding Dehn twists T_1 and T_2 . We have seen that in $\mathbb{P}\mathcal{ML}(\Sigma)$, iterating T_1 and T_2 on \mathcal{L} results in convergence to c_1 and c_2 respectively. Consider $T_1^N(\mathcal{L})$, where N is enormous. The resulting lamination is very close to c_1 . It intersects c_2 , but very little. If we start applying T_2 to $T_1^N(\mathcal{L})$, the lamination starts to converge to c_2 but it does so at a very slow rate. The larger N is, in fact, the slower this convergence is going to be. This is problematic from the point of view of ping-pong, if we want to look at laminations which are close to the limiting laminations of the mapping classes in question as we do in Theorem 6.6, for instance.

It is better to consider laminations which do not necessarily “look like” c_1 or c_2 , but which share some of their properties. A fundamental property we can exploit here is that $i(c_i, c_i) = 0$ for $i = 1, 2$. For $\epsilon > 0$, we will write $U_{i,\epsilon}$ for the set of laminations whose weighted intersection with c_i is at most ϵ .

Suppose that γ is a third curve which intersects both c_1 and c_2 exactly once. Notice that for any choices of n and m , we have that the absolute geometric intersection number

$$i_g(c_j, T_1^n T_2^m(\gamma)) = 1$$

for $j = 1, 2$. Viewing γ as a lamination, we see that as long as the length of $T_1^n T_2^m(\gamma)$ is sufficiently large (in some hyperbolic metric) then the intersection number between the laminations c_j and $T_1^n T_2^m(\gamma)$ is very nearly zero. We can neatly summarize these observations:

Proposition 6.7. *For each N , write A_N for the group $\langle T_1^N, T_2^N \rangle$. For any $\epsilon > 0$, any N sufficiently large, and any $1 \neq g \in A_N$, we have*

$$g(\gamma) \in U_{1,\epsilon} \cap U_{2,\epsilon}.$$

Proof. For simplicity, we will assume that the homology classes of c_1 and c_2 are linearly independent. The technical assumptions on the curves $\{c_1, c_2, \gamma\}$ simplify the argument needed to prove the proposition. In particular, the homology intersection pairing between γ and both c_1 and c_2 is nonzero. In particular, the homology class of $T_1^n T_2^m(\gamma)$ is given by

$$[T_1^n T_2^m(\gamma)] = [\gamma] + n[c_1] + m[c_2].$$

It follows that the word length of $T_1^n T_2^m(\gamma)$ as an element of $\pi_1(\Sigma)$ (which is asymptotically comparable to the hyperbolic length) is at least $n + m$. It follows that the (lamination) intersection number of $T_1^n T_2^m(\gamma)$ with c_1 and c_2 tends to zero at least as fast as $1/(n + m)$, which establishes the claim of the proposition. \square

We will now generalize the preceding proposition and dispense with some of the assumptions. Namely, we suppose that c_1 and c_2 are any two disjoint simple closed curves and γ is just some simple closed curve which intersects both c_1 and c_2 . We will write n_i for the geometric intersection number of γ with c_i , T_i for the Dehn twist about c_i , and $U_{i,\epsilon}$ for the subset of $\mathbb{P}\mathcal{ML}(\Sigma)$ as above. Observe that for any exponents a and b , we have that $T_1^a T_2^b(\gamma)$ intersects c_i no more than n_i times.

Proposition 6.8. *For each N , write A_N for the group $\langle T_1^N, T_2^N \rangle$. For any $\epsilon > 0$, any N sufficiently large, and any $1 \neq g \in A_N$, we have*

$$g(\gamma) \in U_{1,\epsilon} \cap U_{2,\epsilon}.$$

Proof. Consider an intersection point of γ and c_i . Fixing a hyperbolic metric ℓ on Σ , we may represent γ and c_i by geodesics, after which the intersection is necessarily transverse. Lifting γ and c_i to two intersecting geodesics in hyperbolic space, we see that the effect of doing a Dehn twist about c_i is to cut γ at the intersection point, translate one piece along c_i by a distance $\ell(c_i)$, and straighten the resulting union of two rays and a segment to a geodesic. Since c_1 and c_2 are disjoint, we can perform this cutting and straightening construction independently for the two curves, at each intersection point.

A piecewise geodesic lift of $g(\gamma)$ to \mathbb{H}^2 can be given by following γ until we hit an intersection point with a lift of c_i , following c_i for the appropriate distance, and then continuing along another lift of γ .

We claim that there are only a finite number of exponents e_i such that $T_1^{e_1} T_2^{e_2}(\gamma)$ is shorter than any given length. For any compact subset K of \mathbb{H}^2 , there are only finitely many lifts of γ , c_1 and c_2 which intersect K . Furthermore, each lift of c_1 and c_2 to \mathbb{H}^2 cuts \mathbb{H}^2 into two half-spaces. These half-spaces have the property that whenever our piecewise geodesic representative crosses from one such half-space to another, it never backtracks. This can be seen easily from the fact that c_1 and c_2 are disjoint. Finally, each piece of $\gamma \setminus (\gamma \cap (c_1 \cup c_2))$ has length bounded away from zero, since the distance between c_1 and c_2 on Σ is positive. We see that the piecewise geodesic representative segment for $g(\gamma)$ makes positive progress towards infinity as the exponents for T_1 and T_2 tend to infinity. In particular, $g(\gamma)$ can have both of its endpoints in a fixed compact set K for only finitely many different exponents. It follows that choosing N sufficiently large, for any $1 \neq g \in A_N$, we have

$$g(\gamma) \in U_{1,\epsilon} \cap U_{2,\epsilon},$$

since the total number of intersections of $g(\gamma)$ with c_1 and c_2 does not change but since the length of $g(\gamma)$ tends to infinity. \square

The proof of Proposition 6.8 is robust in the sense that it generalizes beyond Dehn twists. Consider, for instance, two pseudo-Anosov homeomorphisms ψ_1 and ψ_2 supported on disjoint surfaces with stable laminations \mathcal{L}_1 and \mathcal{L}_2 , and a simple closed curve γ which intersects both of them essentially. We can analyze the actions of ψ_1 and ψ_2 on γ in a vein which is similar to the proof of Proposition 6.8. One can similarly show that for each n , there are only finitely many exponents e_1 and e_2 for which $\ell(\psi_1^{e_1} \psi_2^{e_2}(\gamma)) \leq n$.

6.5. The issue with boundary twisting. In this subsection and for the rest of this section of the paper, we will fix a hyperbolic metric on the surface Σ , and all laminations under consideration will be assumed to be geodesic. Let ψ be a pseudo-Anosov mapping class supported on a subsurface $\Sigma' \subset \Sigma$, and let β be a boundary component of Σ' . Write \mathcal{L}^\pm for the limiting laminations of ψ . It is possible that the laminations \mathcal{L}^\pm contain β , in which case ψ twists about the boundary component β . We call β a **peripheral leaf** of \mathcal{L}^\pm .

The situation where $\beta \subset \mathcal{L}^\pm$ is pathological for two reasons. Firstly, various angle estimates which we will be making (see Lemma 6.11, for instance) will not work, since β will not be disjoint from \mathcal{L}^\pm . More importantly, Theorem 1.1 fails when we allow peripheral leaves in pseudo-Anosov laminations.

To see this second claim, we will give an example where we get unexpected commutation. Let ψ and ϕ be pseudo-Anosov mapping classes on Σ' which differ by a twist about β , and let φ be another pseudo-Anosov mapping class supported on Σ' which does not commute with ψ or ϕ . Theorem 1.1 says that for sufficiently large N , we have $\langle \psi^N, \phi^N, \varphi^N \rangle \cong \mathbb{Z}^2 * \mathbb{Z}$, where the first factor is

generated by $\{\psi^N, \phi^N\}$. However, $\psi^{-N}\phi^N$ is different from the identity and commutes with φ , a contradiction.

Therefore, we will always assume that pseudo-Anosov laminations on subsurfaces of Σ have no peripheral leaves, or equivalently that each leaf of \mathcal{L}^\pm is dense. A pseudo-Anosov mapping class which twists about a boundary component must be thought of as a pseudo-Anosov mapping class which does not twist about the boundary, composed with a boundary twist.

Theorem 1.1 will still hold if some mapping class ψ on Σ' with boundary twisting is included, provided that the following assumptions hold:

- (1) No pseudo-Anosov ϕ with the same support as ψ virtually commutes with ψ .
- (2) No other mapping class φ is supported on the boundary of Σ' .

6.6. Angles of intersection. In this subsection we will develop some of the necessary technical tools to prove Theorem 1.1. The general strategy for the proof of Theorem 1.1 is to play ping-pong on $\mathbb{P}\mathcal{ML}(\Sigma)$. The fundamental technical difficulty which arises can be illustrated by the following example: let x , y and z be three simple closed curves on Σ with z disjoint from both x and y and with x and y intersecting nontrivially. Suppose γ is a simple closed curve which intersects all three of these curves geometrically plenty of times, but in $\mathbb{P}\mathcal{ML}(\Sigma)$, the lamination γ is very close to z , so that the intersection pairing of z and γ as measured laminations is very small. Now, perform a moderately-sized power of a Dehn twist about y . How can we control the intersection pairing of the resulting curve with z ?

The idea is to control the angles at which these various curves intersect. As observed before, the total number of geometric intersections between γ and the three curves in question cannot increase under performing Dehn twists. If we can guarantee that γ stays relatively parallel to z after Dehn twisting about y , then the lamination intersection pairing between $T_y^N(\gamma)$ and z cannot increase dramatically.

Consider an intersection point between two geodesics in \mathbb{H}^2 . Since the nonempty intersection of any two distinct geodesics in \mathbb{H}^2 is transverse, the intersection has a well-defined angle which is contained in the open interval $(0, \pi)$. Of course, whether two geodesics intersect with angle zero or angle π is a matter of convention which can be established unambiguously by orienting the geodesics. To avoid this problem, we will declare the angle of intersection of two geodesics to be the acute angle at the intersection point (or $\pi/2$ if the two angles are the same). We will call this angle measure the **unoriented angle of intersection**. We will often say that two geodesics on a surface Σ have an unoriented angle of intersection at least $\alpha > 0$. This means that at least one intersection point between those two geodesics has such an unoriented angle of intersection. If we say that the minimal unoriented angle of intersection is at least $\alpha > 0$, we mean that each intersection between those two geodesics has angle at least α .

The following result roughly says that different laminations have no parallel leaves.

Lemma 6.9. *Let γ be a simple closed geodesic on Σ and let \mathcal{L} be the limiting geodesic lamination of a pseudo-Anosov mapping class ψ on a connected subsurface of Σ . There is a $\theta > 0$ such that the minimal unoriented angle of intersection between γ and any leaf of \mathcal{L} is at least θ . The same conclusion holds if γ is replaced by a limiting lamination \mathcal{L}' of another pseudo-Anosov mapping class ψ' , provided that ψ and ψ' do not generate a virtually abelian subgroup of $\text{Mod}(\Sigma)$.*

Proof. When γ is a closed geodesic, this is obvious. If there were a sequence of intersection points of γ with \mathcal{L} with angle tending to zero then γ could not have finite length.

For the case of \mathcal{L}' and \mathcal{L} , consider the a sequence $\{x_n\}$ of intersection points between leaves whose angles tend to zero at least as fast as $\{1/n\}$. Note that \mathcal{L}' and \mathcal{L} are compactly supported. So, the sequence $\{x_n\}$ has an accumulation point x . Both \mathcal{L}' and \mathcal{L} are closed, so that x is in fact an intersection point between a leaf of \mathcal{L} and \mathcal{L}' . It follows that at x , there is a leaf of \mathcal{L} and a leaf

of \mathcal{L}' which coincide. Since each leaf is dense in both laminations, we see that $\mathcal{L} = \mathcal{L}'$. The lemma follows. \square

If \mathcal{L} in Lemma 6.9 is replaced by a closed geodesic, then the same conclusions obviously hold (provided that γ does not coincide with \mathcal{L}). It is clear that the requirement that pseudo-Anosov laminations have no peripheral leaves is essential.

The following result roughly says that applying powers of a mapping class f to a geodesic γ will send intersection angles between γ and limiting laminations of f to zero.

Lemma 6.10. *Fix $\epsilon > 0$, and let \mathcal{L}^\pm be the limiting geodesic laminations of a mapping class f , which is either pseudo-Anosov on a connected subsurface or of a Dehn twist. Let γ be a geodesic which intersects \mathcal{L}^- at least once with unoriented angle at least $\alpha > 0$. Then there is an $N > 0$ which depends only on α , ϵ and f , such that for all $n \geq N$, some intersection between $f^n(\gamma)$ and a leaf of \mathcal{L}^+ has unoriented angle of intersection at most ϵ .*

The distinction between the intersection of γ with \mathcal{L}^- and with \mathcal{L}^+ only applies when f is pseudo-Anosov.

Proof of Lemma 6.10. Lift \mathcal{L} to \mathbb{H}^2 and fix a lift $\tilde{\gamma}$ of γ and a basepoint $x \in \tilde{\gamma}$ which is disjoint from the lift of \mathcal{L} . Follow $\tilde{\gamma}$ from x until we encounter a leaf of the lift of \mathcal{L} . By choosing x carefully, we may assume that the unoriented angle with which we encounter this first leaf is at least α .

Suppose first that f is a Dehn twist, so that $\mathcal{L}^\pm = \mathcal{L}$ is just a simple closed curve on Σ . The effect of the Dehn twist is to cut open $\tilde{\gamma}$, insert a segment along the lift of \mathcal{L} of length $\ell(\mathcal{L})$, and then continue along a translate of $\tilde{\gamma}$. We thus obtain a piecewise geodesic. In order to find the angle between $f(\gamma)$ and \mathcal{L} , we need to straighten out the piecewise geodesic to be geodesic. The conclusion of the lemma now follows from the thinness of triangles in \mathbb{H}^2 . A straightened geodesic representing $f(\gamma)$ must stay within a universally bounded distance of the piecewise geodesic which follows $\tilde{\gamma}$ from x and then a piece of the lift of \mathcal{L} . If the inserted piece of the lift of \mathcal{L} is long enough compared to α , then clearly a geodesic representing $f(\gamma)$ intersects \mathcal{L} at as small an angle as we wish. Thus, fixing ϵ taking N sufficiently large, we can arrange this intersection to have unoriented angle at most ϵ .

Of course, $\tilde{\gamma}$ will generally intersect infinitely many lifts of \mathcal{L} . Consider again $\tilde{\gamma}$ emanating from x and encountering a lift of \mathcal{L} . We insert a segment along that lift of \mathcal{L} and then continue along a translate $\tilde{\gamma}'$ of $\tilde{\gamma}$. Call the direction along which we travel on the lift of \mathcal{L} “positive”, and write p for the corresponding endpoint at infinity in the positive direction.

The translate $\tilde{\gamma}'$ of $\tilde{\gamma}$ defines a geodesic ray emanating from this lift of \mathcal{L} . Fixing x within a fundamental domain for Σ , the point at infinity where this ray hits $\partial\mathbb{H}^2$ depends only on the length of the inserted geodesic segment along the lift of \mathcal{L} and the angle α . By increasing the length of the inserted geodesic segment (which is to say by increasing N), we can avoid any finite collection of lifts of \mathcal{L} in the half-space defined by the one from which the geodesic ray emanates. Let ξ be the next lift of \mathcal{L} which we encounter along $\tilde{\gamma}'$. The endpoint of the geodesic lift of $f(\gamma)$ must be contained between the endpoints of ξ . Since all lifts of \mathcal{L} are disjoint, we can arrange ξ to have endpoints as close to p as we would like by translating further along the first lift of a leaf of \mathcal{L} . Thus, if N is large enough then the other lifts of \mathcal{L} will not undo our estimate that each unoriented angle of intersection between $f^N(\gamma)$ and \mathcal{L} is at most ϵ .

Now suppose that f is pseudo-Anosov on a subsurface. The requirement that γ has unoriented angle of intersection bounded away from zero with \mathcal{L}^- precludes γ from fellow traveling leaves of the unstable lamination of f . The argument is nearly the same as for Dehn twists, only that the length of the twisting curve is replaced by the dilatation of the pseudo-Anosov.

Pull back \mathcal{L}^\pm to \mathbb{H}^2 . Consider a lift $\tilde{\gamma}$ of γ . If we fix a basepoint x and follow $\tilde{\gamma}$, eventually we will reach leaves of the lift of \mathcal{L}^+ which $\tilde{\gamma}$ traverses, and then we will reach a complementary region of \mathcal{L}^+ . The effect of the pseudo-Anosov f on γ is by stretching the leaves of the stable lamination

by the dilatation and contracting the leaves of the unstable lamination by the dilatation. If γ intersects a leaf of the unstable foliation with an unoriented angle of intersection of at least $\alpha > 0$, then f will stretch γ in the direction which is parallel to the leaves of \mathcal{L}^+ at a definite rate which is proportional to the dilatation. The rest of the proof is the same as in the Dehn twist case. \square

Lemma 6.11. *Let \mathcal{L}^\pm be the limiting geodesic laminations of a mapping class f and let X be a limiting geodesic lamination of another mapping class, with X disjoint from \mathcal{L}^\pm . There exists a continuous function*

$$\Theta : (0, \pi/2] \rightarrow (0, \pi/2]$$

which depends only on \mathcal{L}^\pm and X such that:

- (1) *The value of Θ tends to zero as the argument tends to zero.*
- (2) *Let γ be a geodesic which intersects a leaf of X with unoriented angle of intersection $\alpha(\gamma, X)$. Then for any n , there is an intersection between $f^n(\gamma)$ and a leaf of X whose angle of intersection $\alpha(f^n(\gamma), X)$ satisfies*

$$\alpha(f^n(\gamma), X) \leq \Theta(\alpha(\gamma, X)).$$

Proof. Lift all the geodesics in question to \mathbb{H}^2 . Since \mathcal{L}^\pm and X are disjoint, no leaf of \mathcal{L}^\pm can share an endpoint with any leaf of X at infinity. This is clear for any closed leaves of \mathcal{L}^\pm and X since closed leaves represent homotopy classes of simple closed curves and the discreteness of the representation $\pi_1(\Sigma) \rightarrow PSL_2(\mathbb{R})$ precludes disjoint closed geodesics from sharing endpoints at infinity. In general, we can cut Σ open along one of the laminations \mathcal{L}^\pm to get a surface Σ' of the same area as Σ but with geodesic boundary. For the details of this construction, see [49], Proposition 1.6.3. We then double Σ' along this boundary. The resulting surface contains a copy of X since \mathcal{L}^\pm and X are disjoint, and the image of \mathcal{L}^\pm is a geodesic lamination with finitely many leaves. These leaves are all either closed geodesics or geodesics connecting cusps in the doubled surface. Since X is compactly supported on Σ and since no leaf of X spirals into any leaf of \mathcal{L}^\pm , we have that no leaf of X enters any cusps of the doubled surface. Repeating the cutting and doubling construction for X , we get two laminations, each consisting of finitely many leaves, and each leaf connects cusps or is closed. It follows easily now that no two leaves of \mathcal{L}^\pm and X share fixed points at infinity when lifted to \mathbb{H}^2 .

Fix a lift $\tilde{\gamma}$ of γ to \mathbb{H}^2 and a leaf β of the lift of X which $\tilde{\gamma}$ intersects at a point p . We will always assume that this intersection point sits inside of a fixed fundamental domain K for Σ , which without loss of generality is compact. Choosing a different point of intersection between different lifts in a different fundamental domain is tantamount to applying an isometry to the entire picture.

If the angle between $\tilde{\gamma}$ and β is small, then one has to follow $\tilde{\gamma}$ for a very long distance before encountering a leaf of the lift of \mathcal{L}^\pm . Let q_1 be the first leaf of the lift of \mathcal{L}^\pm which $\tilde{\gamma}$ encounters. Observe that for any n , the endpoint at infinity of $f^n(\tilde{\gamma})$ on the side of β which contains q_1 must lie between the endpoints of q_1 . If the unoriented angle of intersection α at p decreases beyond some threshold then $\tilde{\gamma}$ will avoid q_1 altogether, since q_1 and β share no endpoints at infinity. Then, the next candidate q_2 for the first leaf of the lift of \mathcal{L}^\pm which $\tilde{\gamma}$ encounters is a geodesic whose endpoints are wedged between an endpoint of q_1 and an endpoint of β (since $\tilde{\gamma}$ did not encounter q_2 first). Since β and q_2 share no endpoints at infinity, we can repeat the same argument. It follows that as α tends to zero, the set candidates for endpoints of the first leaf of the lift of \mathcal{L}^\pm encountered by $\tilde{\gamma}$ as it travels away from p makes up a smaller and smaller subset of the boundary of \mathbb{H}^2 at infinity, in the visual metric as seen from p .

The observations of the previous paragraph also hold on the other side of β , which is to say as γ moves away from p in the other direction. It follows that for any n , the angle between $f_1^n(\tilde{\gamma})$ and β cannot increase or decrease drastically.

It is clear that the function Θ above can be constructed to be continuous, since it is just measuring the value to which the intersection of $\tilde{\gamma}$ and β can increase as we move endpoints of $\tilde{\gamma}$ between endpoints of lifts of leaves of \mathcal{L}^\pm . \square

6.7. Fellow traveling. Let $\gamma(t)$ and $\gamma'(t)$ be two unit speed parametrized geodesics in \mathbb{H}^2 which intersect at a point p at $t = 0$. It is well-known that as we flow away from time 0, the distance between $\gamma(t)$ and $\gamma'(t)$ diverge quickly. If the angle of intersection at p in the direction of positive time travel is small, then $\gamma(t)$ and $\gamma'(t)$ will stay near each other, at least for a short while.

Lemma 6.12. *Let \mathcal{L} be a simple closed geodesic or a limiting geodesic lamination of a pseudo-Anosov homeomorphism supported on a connected subsurface of Σ . Let X be a geodesic which intersects \mathcal{L} with minimal unoriented angle of intersection $\alpha > 0$. For each $\alpha' < \alpha$, there is a $\theta(\alpha') > 0$ which depends only on \mathcal{L} , X and α' such that if the unoriented angle of intersection between a geodesic γ and \mathcal{L} is at most θ then γ intersects X with unoriented angle at least α' .*

Proof. Lift the entire setup to the universal cover \mathbb{H}^2 . Suppose first that \mathcal{L} is a simple closed curve of length ℓ . Then each component of the lift of \mathcal{L} intersects a component of the lift of X , and the maximum distance between two such intersections along one lift of \mathcal{L} is ℓ . Let a lift $\tilde{\gamma}$ of γ intersect a lift of \mathcal{L} . Traveling along \mathcal{L} , we will encounter a lift of X at an unoriented angle of α at most a distance ℓ away. If the unoriented angle between the lift of \mathcal{L} and $\tilde{\gamma}$ is small enough, $\tilde{\gamma}$ will intersect a lift of X with unoriented angle nearly α .

Now suppose that $\mathcal{L} = \mathcal{L}^+$ is the limiting lamination of a pseudo-Anosov on a subsurface Σ' . Parametrize the leaves of \mathcal{L} to have speed one. Note that if T is a transversal to the lamination then the first return time of leaves of \mathcal{L} to T is bounded. This follows from several well-known facts: the first is that \mathcal{L} can be transformed into a singular foliation \mathcal{F}^+ by collapsing complementary components of \mathcal{L} . Together with the singular foliation \mathcal{F}^- associated to \mathcal{L}^- , these foliations give Σ' a flat metric which is comparable to the hyperbolic metric (cf. [24] for instance). Passing to a double branched cover if necessary, we may assume \mathcal{F} is orientable. Then, the first return map for a transversal is just an interval exchange map and thus has bounded first return time. This establishes the claim.

We now obtain the conclusion of the lemma, since then we can take a neighborhood of an intersection point of X and \mathcal{L} to be a transversal. If γ follows a leaf of \mathcal{L} closely for the maximum first return time to the transversal, then γ will intersect the transversal as well. By taking θ small enough, we can make the unoriented angle of intersection as close to the unoriented angle of intersection between X and \mathcal{L} as we like. \square

6.8. Greedy normal forms in right-angled Artin groups. Before we give the proof of Theorem 1.1, we will describe a normal form which we will use in the course of the proof.

Let $w \in A(\Gamma)$ be a reduced word. We will rewrite w as a product

$$w = w_k w_{k-1} \cdots w_1$$

of central words in a certain way. We will want no generator which occurs in w_i to commute with w_{i+1} . To achieve this, first we require that w_1 be as long as possible, subject to the condition that it still be a central word. Then, we arrange for w_2 to be as long as possible as a subword of $w \cdot w_1^{-1}$, subject to the requirement that it be central. We continue in this fashion until we obtain an expression for w as a product of central words.

Now suppose that a generator v which occurs in w_{k-1} commutes with w_k . We move all occurrences of v to w_k , thus shortening w_{k-1} . Repeat this moving of generators from w_{k-1} to w_k until each generator occurring in w_{k-1} does not commute with w_k . If w is not itself a central word, w_{k-1} will be nontrivial at the end of this process. We now repeat this process for w_{k-2} . If v' is a generator occurring in w_{k-2} which commutes with w_{k-1} then move each occurrence of v' to w_{k-1} . If v' commutes with w_k as well, move the occurrences of v' in w_{k-2} to w_k . In general, if we have

a generator in w_i which commutes with w_{i+1} , we move all occurrences of this generator as far to the left as possible. This rewriting process will terminate in finite time. Note that after we move generators occurring in w_i as far to the left as possible, we may move new generators to w_i from some w_j with $j < i$. However, the generators of w_i which were there already will not be moved to the right. The resulting expression of w as a product of central words will be called **left greedy normal form**. For our purposes, no uniqueness of the left greedy normal form is required – only that it exists.

6.9. Playing ping-pong and the proof of Theorem 1.1. Let $\{f_1, \dots, f_k\}$ be the mapping classes in question. We assume that this collection of mapping classes is irredundant. We write

$$L = \{\mathcal{L}_1^\pm, \dots, \mathcal{L}_k^\pm\}$$

for the invariant laminations of these mapping classes, which are all either pseudo-Anosov laminations supported on connected subsurfaces of Σ in which case \mathcal{L}_i^+ and \mathcal{L}_i^- are different, or simple closed curves, in which case $\mathcal{L}_i^+ = \mathcal{L}_i^-$. A subcollection of L whose elements are all disjoint is called an **admissible configuration**. An admissible configuration \mathcal{C} naturally gives rise to an embedded simplex $\Delta_{\mathcal{C}}$ inside of $\mathbb{P}\mathcal{ML}(\Sigma)$ by taking non-negative total weight one linear combinations of the elements in the subcollection.

Proof of Theorem 1.1. All laminations and curves under consideration are geodesic. For any pair of distinct, intersecting laminations \mathcal{L}_i and \mathcal{L}_j on the list L , write $\epsilon_{i,j}$ for the minimal unoriented angle of intersection between leaves of \mathcal{L}_i and \mathcal{L}_j . Notice that if f_i is pseudo-Anosov on a subsurface, then we include the minimal unoriented angle of intersection between \mathcal{L}_i^+ and \mathcal{L}_i^- . Let γ be a simple closed curve which intersects each element of L essentially (and is therefore not on the list L). Since L is a finite set, there is a positive lower bound ϵ_γ on the minimal unoriented angle of intersection between γ and leaves of elements of L . Write

$$\delta = \frac{\min(\epsilon_{i,j}, \epsilon_\gamma)}{2}.$$

The value of δ is positive by Lemma 6.9.

Let $\epsilon > 0$, which we will assume for the rest of the proof is smaller than δ . By Lemma 6.10, there is an exponent $N = N(\delta) > 0$ such that if γ is any geodesic which intersects the limiting laminations \mathcal{L}^\pm of some f_i with unoriented angle of intersection at least δ , then some unoriented angle of intersection between \mathcal{L}_i^+ and $f_i^n(\gamma)$ is at most ϵ for all $n \geq N$. Suppose that instead of $f_i^N(\gamma)$ we consider

$$f_i^n g_1^{n_1} \cdots g_j^{n_j}(\gamma),$$

where each one of $\{g_1, \dots, g_j\}$ is equal to a mapping class on the list $\{f_1, \dots, f_k\}$ and all the $\{g_1, \dots, g_j\}$ commute with each other and with f_i . In this case, Lemma 6.11 gives us a function Θ such that the unoriented angle of intersection between

$$f_i^n g_1^{n_1} \cdots g_j^{n_j}(\gamma)$$

and \mathcal{L}_i^+ will be at most $\Theta^j(\epsilon)$, where by Θ^j we mean Θ composed with itself j times. We take Θ here to be the minimum of the functions furnished by Lemma 6.11 which work for pairs of disjoint laminations in L .

For any fixed $\epsilon > 0$ and j , and for all sufficiently small ϵ_0 , we have $\Theta^j(\epsilon_0) < \epsilon$. The maximal size of a commuting collection of mapping classes on the list $\{f_1, \dots, f_k\}$ is bounded, so that j is always bounded. Therefore, given $\epsilon_0 > 0$, we can choose N sufficiently large so that if $\{f_i, g_1, \dots, g_j\}$ are a commuting collection in $\{f_1, \dots, f_k\}$, we have that the unoriented angle of intersection between \mathcal{L}_i and

$$f_i^n g_1^{n_1} \cdots g_j^{n_j}(\gamma)$$

is at most ϵ , whenever $n \geq N$.

In the first paragraph, we fixed a simple closed curve γ which intersects each element of L . By replacing the $N(\delta)$ in the previous paragraph by a possibly larger exponent, we can arrange that the unoriented angle of intersection between $f_i^n(\gamma)$ and \mathcal{L}_i^\pm is at most ϵ_0 .

Recall that in Lemma 6.12, we showed that if a geodesic γ follows some \mathcal{L} closely enough then γ will intersect any other \mathcal{L}' which \mathcal{L} intersects, with unoriented angle close to that of \mathcal{L} and \mathcal{L}' . For each pair of laminations associated to mapping classes on the list $\{f_1, \dots, f_k\}$, there is a $\theta_0 > 0$ such that if γ intersects \mathcal{L} with unoriented angle at most θ_0 then γ will intersect \mathcal{L}' with unoriented angle at least δ . We will be insisting that $\epsilon < \theta_0$. Henceforth in the proof, fix ϵ and ϵ_0 which satisfy the conditions we have specified above.

Let Γ be the coincidence graph of the limiting laminations of elements of L , so that there is a vertex for each pair \mathcal{L}_i^\pm and an edge between \mathcal{L}_i^\pm and \mathcal{L}_j^\pm if and only if \mathcal{L}_i^\pm and \mathcal{L}_j^\pm are disjoint. Labeling the vertices of Γ by $\{v_1, \dots, v_n\}$, for each N we obtain a homomorphism

$$\phi_N : A(\Gamma) \rightarrow \text{Mod}_{g,p}$$

which sends v_i to f_i^N . We will show that for all N sufficiently large, this homomorphism is injective. For each N , we will not distinguish notationally between elements of $A(\Gamma)$ and their image in $\text{Mod}_{g,p}$ under ϕ_N . The particular N we will be choosing is the one which makes Lemma 6.10 work for $\{f_1, \dots, f_k\}$, with $\alpha = \delta$ and ϵ_0 as above.

Let $U_{i,\epsilon}^+, U_{i,\epsilon}^- \subset \mathbb{P}\mathcal{ML}(\Sigma)$ be the laminations on Σ such that if $c \in U_{i,\epsilon}^+$ then there is a leaf of c which intersects a leaf of \mathcal{L}_i^+ with unoriented angle at most ϵ , and respectively for $U_{i,\epsilon}^-$ and \mathcal{L}_i^- . For each admissible configuration \mathcal{C} , we will write

$$U_{\mathcal{C},\epsilon} = \bigcap_{\mathcal{L} \in \mathcal{C}} U_{\mathcal{L},\epsilon},$$

where $U_{\mathcal{L},\epsilon} = U_{i,\epsilon}^\pm$, depending on whether $\mathcal{L} = \mathcal{L}_i^+$ or \mathcal{L}_i^- .

We can now finish the proof of the result. The proof will use a version of ping-pong and the expression of each word $w \in A(\Gamma)$ in left greedy normal form. For the base case, let w be a central word. Write \mathcal{C} for the admissible configuration which is the support of w , viewed as a mapping class. Let α be the minimal unoriented angle of intersection between γ and \mathcal{L}_i^\pm , as i is allowed to vary. Since ϵ was smaller than δ and since $\alpha > \delta$, we have that

$$\gamma \notin \bigcup_i U_{i,\epsilon}^\pm$$

for any i .

We have already shown in the second paragraph that $w(\gamma)$ has unoriented angle of intersection at most ϵ with each component of \mathcal{C} . In particular, $w(\gamma) \in U_{\mathcal{C},\epsilon}$.

Now suppose $c \in U_{i,\epsilon}^\pm$ and that f_j does not commute with f_i . By Lemma 6.12, we may assume that c intersects \mathcal{L}_j^\pm with angle at least δ . In particular, $f_j^n(c) \in U_{j,\epsilon_0}^+$ for all $n \geq N$ (and similarly $f_j^{-n}(c) \in U_{j,\epsilon_0}^-$ for all $n \geq N$).

We claim that this proves the result. Indeed, write $w = w_k \cdots w_1$ in left greedy normal form. We have that $w_1(\gamma) \in U_{i,\epsilon}^\pm$ for some i . We may assume that $w_{k-1} \cdots w_1(\gamma) \in U_{j,\epsilon}^\pm$, that that there is a vertex generator occurring in w_k , corresponding to the mapping class f_m , which does not commute with f_j . Any curve in $U_{j,\epsilon}^\pm$ intersects \mathcal{L}_m^\pm with unoriented angle at least δ , so that combining Lemmas 6.10 and 6.11, we see that for any $n \geq N$,

$$f_m^n(w_{k-1} \cdots w_1(\gamma)) \in U_{m,\epsilon_0}^+,$$

so that

$$w(\gamma) = w_k(w_{k-1} \cdots w_1(\gamma)) \in U_{m,\epsilon}^+,$$

and similarly after replacing f_m^n by a negative power of f_m , and replacing $U_{m,\epsilon}^+$ by $U_{m,\epsilon}^-$. In particular, $w(\gamma) \neq \gamma$. \square

7. RIGHT-ANGLED ARTIN GROUPS IN $\text{Mod}_{g,p}$ AND THE ISOMORPHISM PROBLEM FOR SUBGROUPS OF $\text{Mod}_{g,p}$

Recall the precise formulation of the isomorphism problem for a class of finite presented groups (see Miller’s book [46] for a classical account of decision problems in group theory): We are given a recursive class $\Phi = \{\Pi_i\}$ of finitely presented groups, and we wish to find an algorithm ϕ which on the input (i, j) determines whether or not there exists an isomorphism between Π_i and Π_j . It is an open problem to solve the isomorphism problem for finitely presented subgroups of $\text{Mod}_{g,p}$ (see Farb’s article [21]). Usually, it is required that Φ be a class of finitely presented groups, though in our setup it will be more natural to require that they be merely finitely generated in our setup.

7.1. Decision problems in $\text{Mod}_{g,p}$. Let $A \subset \text{Mod}_{g,p}$ be a finite generating set. We may recursively enumerate words in the free group $F(A)$ as $\{a_1, a_2, \dots\}$, by choosing an ordering on $A \cup A^{-1}$ and listing words in lexicographical order and by increasing length, for instance. Since $\text{Mod}_{g,p}$ has a solvable word problem with respect to any finite generating set, we may omit any a_i which is equal in $\text{Mod}_{g,p}$ to some a_j with $j < i$. We may thus recursively enumerate all finitely generated subgroups of $\text{Mod}_{g,p}$: for each n , enumerate all subsets of $A_n = \{a_1, \dots, a_n\}$, omitting those which were enumerated for some $m < n$. We then associate to each finite subset of A_n the group generated by those elements.

Proposition 7.1. *There is no algorithm which determines whether or not two finitely generated subgroups of $\text{Mod}_{g,p}$ are isomorphic.*

This proposition was known by the work of Stillwell in [54] in the context of the **occurrence problem**, also known as the **membership problem**. The membership problem was originally formulated by Mihailova in [44]. The membership problem seeks to determine whether a given element $w \in G$ is contained in a subgroup $\langle F \rangle < G$ generated by a finite set F . Related discussion can be found in Hamenstädt’s article [31].

Proposition 7.1 follows from a more general result about groups which contain sufficiently complicated products of free groups. If G is any finitely generated group, we can recursively enumerate the elements of G , and hence all finitely generated subgroups of G as for $\text{Mod}_{g,p}$.

We will be relying fundamentally on the following well-known result (see for instance Lyndon and Schupp’s book [43], or Miller’s book [46]):

Lemma 7.2. *$G = F_n \times F_n$ has an unsolvable generation problem whenever $n \geq 6$. Precisely, it is undecidable whether or not a finite subset $F \subset G$ generates all of G .*

The following result was claimed in the introduction:

Lemma 7.3. *There is no algorithm which determines whether two finitely generated subgroups of $F_2 \times F_2$ are isomorphic.*

Lemma 7.3 is well-known, and a proof can be found in [47], for instance. It is a theorem of Bridson and Miller (see [9]) that there is no algorithm to determine the homology of an arbitrary finitely generated subgroup of $F_2 \times F_2$. On the other hand, they prove that the isomorphism problem for finitely presented subgroups is solvable. We remark that G. Baumslag and J.E. Roseblade showed in a precise sense that “most” finitely generated subgroups of a product of two nonabelian free groups are not finitely presented (see [2]). On a related theme, G. Levitt has proven the unsolvability of the isomorphism problem for free abelian-by-free groups (see [40]).

Proof of Lemma 7.3. Whenever we produce a generating set for $F_2 \times F_2$, we will suppose that it consists of a generating set for $F_2 \times \{1\}$ and a generating set for $\{1\} \times F_2$. Once we have such a generating set, it will be easy to produce a generating set for $F_6 \times F_6$ sitting inside of $F_2 \times F_2$, using the Nielsen–Schreier rewriting process (see [43]).

Suppose the existence of an algorithm which determines whether two finitely generated subgroups of $F_2 \times F_2$ are isomorphic. Using Stallings' folding automaton (see [53]), we can determine whether a given finite subset of $F_2 \times F_2$ with respect to our nice generating set sits inside of $F_6 \times F_6$. We will show that this algorithm determines whether a finite subset of $F_6 \times F_6$ generates all of $F_6 \times F_6$, whence the conclusion of the lemma. Given a finite subset S of $F_6 \times F_6$, the algorithm returns whether or not $\langle S \rangle < F_6 \times F_6$ is isomorphic to $F_6 \times F_6$. If the algorithm returns “no” then S clearly does not generate all of $F_6 \times F_6$.

Suppose that the algorithm returns “yes”, so that we have $G = \langle S \rangle \cong F_6 \times F_6$. Decompose G as $G_1 \times G_2$, where each $G_i \cong F_6$. Let p_1 and p_2 denote the projections of $F_6 \times F_6$ onto the two factors. Since G is included into $F_6 \times F_6$, we may suppose that $p_1(G_1)$ is nonabelian. Note that centralizers of nonidentity elements inside a nonabelian free group are cyclic. Therefore if $p_1(G_1)$ is nonabelian, we have $p_1(G_2)$ is trivial. It follows that $p_2(G_2) \cong G_2$, so that $p_2(G_1)$ is trivial. In particular, $p_1(G_1) \cong G_1$.

Thus, we see that S generates all of $F_6 \times F_6$ if and only if the inclusions of $p_i(G_i)$ (which we just denote by G_i) into the factors are surjective. Elements of S are written in terms of generators for $F_6 \times \{1\}$ and $\{1\} \times F_6$, so that G_i is obtained simply by deleting the generators for one of the factors. We are thus given a finitely generated subgroup of F_6 and asked to determine whether it is all of F_6 . It is well-known that this problem is effectively solvable (again by Stallings' folding automaton, for instance). It follows that we can algorithmically determine whether or not S generates a proper subgroup of $F_6 \times F_6$, a contradiction. \square

The proof of Proposition 1.9 is almost immediate now.

Proof of Proposition 1.9. Suppose that we are given G and a copy $F_2 \times F_2 < G$. Restricting the isomorphism problem to finitely generated subgroups of the copy of $F_2 \times F_2$, we see that if we could solve the isomorphism problem for finitely generated subgroups of G then we would be able to solve it for $F_2 \times F_2$, a contradiction of Lemma 7.3. \square

Note that in the proof of Proposition 1.9 we are assuming that we are handed a copy of $F_2 \times F_2$ inside of G . This does not detract from the generality of the discussion. Indeed, suppose we knew that G contains a copy of $F_2 \times F_2$ but we do not know how to immediately identify it. If the claim “the isomorphism problem for finitely generated subgroups of G is unsolvable” were to be contested, then one would need to produce an algorithm which solves the problem. But then one could enumerate finitely generated subgroups of G and find a copy of $F_2 \times F_2$ in finite time. Then, one would be reduced to the situation of having a group G and a copy of $F_2 \times F_2$ identified inside of G .

Proof of Proposition 7.1. By Lemma 7.3, we only need to show that we can effectively identify copies of $F_2 \times F_2$ inside of $\text{Mod}_{g,p}$. Clearly if we have four distinct, essential simple closed curves $\{c_1, \dots, c_4\}$ in Σ such that $\{c_1, c_2\}$ and $\{c_3, c_4\}$ are pairwise disjoint and within both $\{c_1, c_2\}$ and $\{c_3, c_4\}$ we have intersections, then the squares of Dehn twists about these curves will generate a copy of $F_2 \times F_2$. Most presentations of the mapping class group are given in terms of Dehn twists about simple closed curves, so it is usually easy to find such a copy of $F_2 \times F_2$.

If one is given a more complicated generating set, one can add a collection of twists about four curves in the configuration described above into the generating set. It can be shown that the solution to the isomorphism problem for finitely generated subgroups does not depend on the given presentation of the ambient group. To see this easily in our case, the paragraph following the proof of Proposition 1.9 illustrates that if a group G even abstractly contains a group which provides a certificate of the unsolvability of the isomorphism problem for finitely generated subgroups (such as $F_2 \times F_2$) then changing the presentation of G will not change the solvability of the isomorphism problem. \square

7.2. Homological rigidity of right-angled Artin groups: no accidental isomorphisms.

Let Γ and Γ' be two finite graphs. If $\Gamma \neq \Gamma'$ then the standard presentations of the associated right-angled Artin groups $A(\Gamma)$ and $A(\Gamma')$ are different, but it is not obvious that $A(\Gamma)$ and $A(\Gamma')$ are not isomorphic. We now prove the following rigidity theorem, which effectively solves the isomorphism problem for right-angled Artin groups.

The following is Theorem 1.10 from the introduction.

Theorem 7.4. *Let G and G' be two right-angled Artin groups. Then $G \cong G'$ if and only if $H^*(G, \mathbb{Q}) \cong H^*(G', \mathbb{Q})$ as algebras.*

Here we drop the notation “ $A(\Gamma)$ ” for a right-angled Artin group in order to emphasize the fact that we do not know the underlying graph.

We will show that if we know the cup product structure on the cohomology algebra of G , then we can reconstruct a unique graph Γ for which $G \cong A(\Gamma)$. We will actually prove a slightly stronger statement: Γ is determined by $H^1(G, \mathbb{Q})$, $H^2(G, \mathbb{Q})$, and the cup product

$$\cup : H^1(G, \mathbb{Q}) \otimes H^1(G, \mathbb{Q}) \rightarrow H^2(G, \mathbb{Q}).$$

It is not surprising that such a strengthening holds since a right-angled Artin group is determined only by a graph, not by an entire flag complex.

Thus, if we are given the data of the abstract isomorphism type of a right-angled Artin group G which is secretly the right-angled Artin group (G, S) , we can use the cohomology algebra of G to very nearly recover S , and indeed nearly enough to recover the underlying graph. The isomorphism type of the graph is uniquely determined, so that we recover the main result of [20]. In the end, we will show that if (G, S) is a right-angled Artin group and G' is abstractly isomorphic to G then there is an essentially unique way to find a right-angled Artin system S' such that we obtain an isomorphism of pairs $(G, S) \cong (G', S')$.

Proof of Theorem 1.10. Throughout the proof we will be tacitly identifying vertices of a graph Γ with their Poincaré duals in $H^1(A(\Gamma), \mathbb{Z})$. Note that Poincaré duals of vertex generators do make sense since each clique corresponds canonically to a torus within the Salvetti complex.

Suppose that Γ is a finite graph with vertex set V which is a candidate for a graph which gives rise to G . We claim that the structure of the cohomology algebra allows us extract a basis X for $H_1(G, \mathbb{Q})$ such that there exists a map of sets $X \rightarrow V$ which induces an isomorphism of algebras $H^*(G, \mathbb{Q}) \rightarrow H^*(A(\Gamma), \mathbb{Q})$.

There is a (possibly trivial) subspace $W_0 \subset H^1(G, \mathbb{Q})$ which is an isotropic subspace for the cup product. Its dimension is precisely the number of isolated vertices of Γ . Choose an arbitrary basis X_0 for W_0 . The cup product descends to the quotient space $A = H^1(G, \mathbb{Q})/W_0$.

Consider the vectors $\{v\}$ in A such that the map

$$f_v : A \rightarrow H^2(G, \mathbb{Q})$$

given by $f_v(c) = v \cup c$ has rank exactly one. We think of v as a weighted sum of vertices of Γ . Clearly v cannot be supported on a vertex of degree more than one, since otherwise the rank of f_v would be too large. Suppose that v is supported on at least two vertices of degree one. Then all these vertices must in fact be connected to one single vertex, by rank considerations. Precisely, if $\{v_1, \dots, v_k\}$ is the support of v , then for all i, j , we have $\text{lk}(v_i) = \text{lk}(v_j) = v_{k+1}$ for some other vertex v_{k+1} which is not contained in the support of v .

Choose a basis B_1 for the span of the cohomology classes v such that f_v has rank one, and suppose that each element b in this basis satisfies $\text{rk } f_b = 1$. We may first choose basis vectors which cup trivially with all the other rank one cohomology classes, and these correspond to degree one vertices which are not connected to any other degree one vertices. Since the remaining degree one vertices pair off as disconnected components of the underlying graph, the choice of the other basis vectors is largely irrelevant.

To interpret the rank of f_v properly, suppose that v is supported on vertices v_1, \dots, v_m . Note first that the rank of f_v is at most the number of edges in the union of the stars of v_1, \dots, v_m , but we shall see that in general there is an inequality.

Suppose v is a cohomology class with $\text{rk } f_v = n > 1$. Clearly v is not supported on any vertex of degree greater than n . Suppose that v is supported on multiple vertices of degree n . If v_1 and v_2 are two such vertices, we have n edges emanating out of each of v_1 and v_2 . Note that there can be an overlap of at most one edge between these two collections of edges, since loops of length 2 are not allowed. Thus there are at least $2n - 1$ distinct edges emanating from v_1 and v_2 . Let w_1, \dots, w_n be the vertices adjacent to v_1 , and let z_1, \dots, z_n be the vertices adjacent to v_2 . Note that if there is a vertex z which is adjacent to v_2 but not in the star of v_1 then the rank of f_v is already too large. It follows that

$$\text{st}(v_1) \cup v_2 = \text{st}(v_2) \cup v_1.$$

In this case we say that v_1 and v_2 have **essentially the same star**.

Suppose that $\{v_1, \dots, v_k\}$ are classes in $H^1(G, \mathbb{Q})$ whose ranks (under the map $v \mapsto f_v$) are exactly n , and assume that each v_i is supported on at least one vertex of rank n . Suppose furthermore that every nonzero linear combination of $\{v_1, \dots, v_k\}$ also has rank exactly n and that the total linear span of $\{v_1, \dots, v_k\}$ has dimension k . It follows that from each pair $\{v_i, v_j\}$ in the collection of classes, we can produce a cohomology class which has rank n and is supported on two distinct vertices of degree n . It follows by the observation above that all the degree n vertices which are in the support of $\{v_1, \dots, v_k\}$ have essentially the same star.

Inductively, choose a basis B_i for the span of cohomology classes v for which $\text{rk } f_v = i$, for each i . We start first with rank one classes. Suppose we have chosen a basis for the rank n classes. We choose a linearly independent set of rank $n + 1$ classes. If a given rank $n + 1$ class is not in the span W_i of

$$\bigcup_{i \leq n} B_i,$$

we first see if it is contained in a subspace $V \subset H^1(G, \mathbb{Q})$ of dimension at least 2 and consisting of rank $n + 1$ classes. If a given class is contained in such a subspace V and if $V \cap W_i = \{0\}$, we see that there are at least two vertices with essentially the same star. If we can find such a subspace V , we can choose a basis for it arbitrarily. If no such subspace exists for a given rank $n + 1$ cohomology class v , then v is supported on a unique vertex w of degree $n + 1$ and possibly also on some degree one vertices adjacent to w . In this sense, we can say that w is determined up to **degree one vertex ambiguity**.

It is now easy to reconstruct the graph underlying G . The basis vectors we chose either correspond to unique vertices up to degree one ambiguity, or they naturally come in bunches which correspond to vertices with essentially the same star. Edges of the graph can be reconstructed by cupping basis vectors together. \square

Corollary 7.5. *$A(\Gamma) \cong A(\Gamma')$ if and only if $\Gamma \cong \Gamma'$.*

If G is a group, we say that G has a solvable **generalized isomorphism problem** for finitely generated subgroups if there is an algorithm A and an irredundant list L such that for every finite $\{f_1, \dots, f_n\} = F \subset G$ there is an $N = N(F)$ such that A takes as input the data of F and returns an element of L which is the isomorphism type of $\{f_1^N, \dots, f_n^N\}$. By Theorem 1.1, we thus obtain the following solution to the generalized isomorphism problem for finitely generated subgroups of the mapping class group:

Corollary 7.6. *Let $\{f_1, \dots, f_k\} \subset \text{Mod}_{g,p}$ be powers of Dehn twists about simple closed curves or pseudo-Anosov on connected subsurfaces. Then there exist an N such that for all $n \geq N$, the isomorphism type of $\{f_1^n, \dots, f_k^n\}$ is determined by the intersection pattern of the canonical*

reduction systems $\{\mathcal{C}_i\}$ for each f_i , together with the data of whether f_i and f_j virtually commute on each component of \mathcal{C}_i , \mathcal{C}_j , $\Sigma \setminus \mathcal{C}_i$, and $\Sigma \setminus \mathcal{C}_j$.

In particular, even though the isomorphism problem for finitely generated subgroups of the mapping class is unsolvable, it becomes solvable for a large class of subgroups if we are willing to replace elements of a finite generating set by sufficiently large positive powers. For arbitrary finitely generated subgroups, we would need the following to hold in order to determine the isomorphism type of the subgroup generated by large powers of the corresponding generators:

Conjecture 7.7. *The isomorphism problem for groups enveloped by right-angled Artin groups is solvable.*

We briefly remark that one subtlety in Theorem 1.1 was that we insisted on irredundancy of the collection $\{f_1, \dots, f_k\}$. This does not pose great difficulty for us for the following reason. Suppose that two elements f_i, f_j on our list of mapping classes were equal. Then we could algorithmically determine whether or not they are equal, since the mapping class group has a solvable word problem.

In the setup of Theorem 1.1, irredundancy is equivalent to $\langle f_i, f_j \rangle \neq \mathbb{Z}$ whenever $i \neq j$, which is somewhat more subtle question than just equality. To determine this purely from the algebra of $\text{Mod}_{g,p}$, one needs to show the existence of a mapping class g which commutes with f_i but not f_j , which is nontrivial from a decision problem standpoint. Since we are already insisting that the mapping classes we consider be of a particular geometric form, we simply require that our list L be irredundant.

7.3. Embedding right-angled Artin groups in mapping class groups. We now turn our attention to the question of which right-angled Artin groups appear in a given mapping class group. Such problems have been discussed in [16], for example.

Question 7.8. *Let Γ be a finite graph with vertex set V , $A(\Gamma)$ the associated right-angled Artin group, and let $\text{Mod}_{g,p}$ be a mapping class group. Under what conditions is there an embedding $A(\Gamma) \rightarrow \text{Mod}_{g,p}$? What types of mapping classes can we arrange in the image of V ?*

There are some easy preliminary observations we can make:

Lemma 7.9. *Let Γ be a finite graph with vertex set V , let $v \in V$, and let $k \neq 0$. Then the subgroup of $A(\Gamma)$ generated by v^k and $V \setminus v$ is isomorphic to $A(\Gamma)$.*

Proof. Let X be a set on which $A(\Gamma)$ plays ping-pong, together with subsets $\{X_i\}$ and a basepoint x_0 . Clearly the subgroup generated by $V \setminus v$ is a right-angled Artin group. If there is an edge between $v' \in V \setminus v$ and v then v^k preserves $X_{v'}$. If there is no edge then v^k sends $X_{v'}$ to X_v . Finally, v^k sends x_0 into X_v . \square

In particular, the isomorphism classes of right-angled Artin groups which can be embedded into $\text{Mod}_{g,p}$ are precisely those which arise from replacing a finite collection of mapping classes by sufficiently high powers and looking at the resulting right-angled Artin group. Our observations concerning ping-pong on $\mathbb{PML}(\Sigma)$ show:

Corollary 7.10. *If $\{f_1, \dots, f_k\}$ are pseudo-Anosov mapping classes which generate a right-angled Artin group, then this group is free.*

We first prove two stable, easier embedding theorems for right-angled Artin groups in Σ_g :

Proposition 7.11. *Let Γ be a finite graph and $A(\Gamma)$ the associated right-angled Artin group. Then there exists a closed surface $\Sigma = \Sigma_{g,p}$ such that $A(\Gamma) < \text{Mod}_{g,p}$, and each vertex generator of $A(\Gamma)$ sits in $\text{Mod}_{g,p}$ as a power of a Dehn twist about simple closed curve.*

Proof. Fix a surface Σ of genus $g \geq 1$. Since the curve complex of Σ has infinite diameter, we may choose a collection of simple closed curves on Σ , one for each vertex of Γ , each of which intersects every other. In general position, each intersection is a double intersection, so that there are no triple points. Label the curves by the vertices of Γ . If two curves c_1, c_2 correspond to vertices in Γ which are connected by an edge then we perform the following modification of Σ :

- (1) Find the intersection points of c_1 and c_2 and remove small neighborhoods of the intersection points.
- (2) Glue in a torus with one boundary component ∂ in place of each intersection.
- (3) Complete c_1 and c_2 to simple closed curves through each of these tori in such a way that they do not intersect. This can be done in such a way that the homology classes of the intersections of c_1 and c_2 relative to ∂ are distinct.

There are two facts to be verified now. First that none of the curves in the modified surface Σ' are inessential or pairwise isotopic, and secondly that any two curves whose pairwise intersections were not modified still have positive geometric intersection number. Upon verifying these facts we apply Theorem 1.1 to get the proposition.

Suppose that c is a curve in Σ' . Then c has not become inessential. Indeed, since c itself was not inessential, $\Sigma \setminus c$ has at most two components, each with nontrivial topology. The modifications done to obtain Σ' do one of two things:

- (1) Combine two components of $\Sigma \setminus c$ into one component when c itself is modified.
- (2) Increase the topological complexity of a component of $\Sigma \setminus c$.

Thus, c does not become inessential.

If c_1, c_2 are two curves which intersect in two points, c_1 and c_2 can be deformed by isotopy in order to intersect fewer times if and only if c_1 and c_2 cobound a bigon. If they cobound a bigon in Σ' then one can show that they cobound a bigon in Σ . Finally, c_1 and c_2 are isotopic if and only if they cobound an annulus. If they cobound an annulus in Σ' then one can easily check that they already cobound an annulus in Σ . \square

Proposition 7.12. *Let Γ be a finite graph and $A(\Gamma)$ the associated right-angled Artin group. Then there exists a surface $\Sigma = \Sigma_{g,p}$ such that $A(\Gamma) < \text{Mod}_{g,p}$, and each vertex generator of $A(\Gamma)$ sits in $\text{Mod}_{g,p}$ as a pseudo-Anosov homeomorphism on a subsurface of Σ .*

Proof. Let Γ^* denote the complement of Γ , which is obtained by taking the complete graph on the vertices of Γ and deleting the edges which lie in Γ itself. Find a closed surface Σ_0 of some genus large enough to accommodate Γ^* . For each vertex and each edge, add a genus one handle to Σ_0 to get a surface Σ . For each vertex $v \in \Gamma^*$, let Σ_v be the connected subsurface with one boundary component which contains the handles corresponding to v and all edges emanating from v . If v and w are not adjacent, we arrange so that Σ_v and Σ_w are disjoint. We then let ψ_v be a pseudo-Anosov homeomorphism supported on Σ_v . The conclusion of the proposition follows easily. \square

One can easily modify the previous proof to arrange for Σ to be closed. The following is now immediate:

Corollary 7.13. *For each graph Γ , there is an embedding $A(\Gamma) \rightarrow \text{Mod}_{g,p}$, where the image of the embedding is contained in the Torelli subgroup of $\text{Mod}_{g,p}$.*

Proof. In the proof of the previous proposition, each Σ_v can be arranged to have genus at least two, in which case we can choose ψ_v to sit in the Torelli group. \square

The following corollary was suggested in the introduction, and it is trivial to establish now.

Corollary 7.14. *Let G be enveloped by a right-angled Artin group. Then there is a surface Σ such that $G < \text{Mod}(\Sigma)$.*

We now wish to fix a genus and embed various right-angled Artin groups into $\text{Mod}_{g,p}$. Note that by the main result of [6], there is a rank obstruction on embedding a given right-angled Artin group in a given mapping class group, a fact which is reflected in the structure of the curve complex. This is an example of a local obstruction to embedding a right-angled Artin group, since the obstruction can be determined by understanding the neighborhood of radius 1 about each vertex in Γ .

We will now give a slightly less obvious example of a local obstruction to embedability of right-angled Artin groups in mapping class groups. Let F be a nonabelian free group of finite rank. Note that it is easy to embed a copy of $F \times F$ inside of Mod_g for $g \geq 2$. Now let $P_n = F \times \cdots \times F$ be a finite product of free groups, which we realize as a right-angled Artin group on a graph.

Proposition 7.15. *Suppose $P_n < \text{Mod}_g$. Then $n \leq 3(g-1)/2$ when g is odd and $n \leq 3g/2$ when g is even. These bounds are sharp.*

Proof. First suppose that P_n is generated by pseudo-Anosov homeomorphisms supported on subsurfaces, and that these are the vertex generators of P_n . A subsurface which supports a pseudo-Anosov homeomorphism must either be a torus with one boundary component or have Euler characteristic at most -2 . To realize P_n in the way we wish, we must partition Σ into n subsurfaces of this type. The Euler characteristic of Σ_g is $2 - 2g$, so that if we have n_T tori and n_S other surfaces, we have

$$n_S \leq \frac{2g - 2 - n_T}{2}.$$

Since there are at most g tori, $g/2 - 1 \leq n_S \leq g - 1$. Every time we increase the number of tori by one, n_S decreases by $1/2$. The bound follows for pseudo-Anosov homeomorphisms supported on subsurfaces.

Next, suppose that P_n is generated by twists about multicurves, and that these are the vertex generators of P_n . Choose a multicurve out of each factor of P_n . Clearly there are no more than $3g - 3$ of these multitwists, since we may combine them all and extend to a pants decomposition of Σ_g . Again, the least complicated surfaces which support two non-commuting Dehn twists are the once punctured torus and the four-times punctured sphere. Since factors of P_n commute with each other, we may separate them from each other, in the sense that all the curves which generate twists in different factors are supported on disjoint subsurfaces. The same combinatorial argument as above gives the upper bound.

In the general case, we may assume that each vertex generator of P_n is either a pseudo-Anosov homeomorphism supported on a connected subsurface or a Dehn twist about a simple closed curve, since otherwise we can extract smaller subsurfaces supporting direct factors of a product of free groups P_n . Two non-commuting vertex generators have a pair of non-commuting twists as factors, or one of generators has a pseudo-Anosov component. By the same combinatorial arguments as above, we must be able to build Σ_g out of n once punctured tori and other surfaces of Euler characteristic no more than -2 , so that the desired bound holds.

The sharpness of the bound can be achieved easily by decomposing the surface into g tori with one boundary component and $g/2$ four-times punctured spheres, when g is even. When g is odd, the right bound is $(g-1)/2$ four-times punctured spheres. \square

From the work of [16] it seems likely that there are more global obstructions to embedding particular right-angled Artin groups into mapping class groups, and this is a potential direction for further inquiry. We can say a few basic things:

Recall the notion of the **curve complex** $\mathcal{C}(\Sigma)$. This is the a simplicial flag complex whose vertices are isotopy classes of essential, nonperipheral, simple closed curves in Σ and whose incidence relation is given by simultaneous disjoint realization. It is easy to see that $\mathcal{C}(\Sigma)$ is locally infinite and finite dimensional. We will mostly be concerned with the 1-skeleton of $\mathcal{C}(\Sigma)$, which we also denote by $\mathcal{C}(\Sigma)$.

Let Γ, Δ be graphs and let $f : \Gamma \rightarrow \Delta$ be an embedding. We call f a **flag embedding** if when two vertices of Γ are sent to adjacent vertices of Δ , they were adjacent in Γ . Equivalently, f preserves the adjacency relation, or equivalently $f(\Delta)$ is an induced subgraph of Γ which is isomorphic to Δ .

Proposition 7.16. *Let $A(\Gamma)$ be a right-angled Artin group. Then $A(\Gamma)$ can be embedded in $\text{Mod}_{g,p}$ in such a way that the right-angled Artin system determined by Γ is mapped to powers of twists about simple closed curves if and only if there exists a flag embedding $f : \Gamma \rightarrow \mathcal{C}(\Sigma)$.*

Proof. Suppose that $f : \Gamma \rightarrow \mathcal{C}(\Sigma)$ is a flag embedding. The images of two vertices of Γ are adjacent if and only if their images were adjacent in $\mathcal{C}(\Sigma)$. It is well-known that two twists commute if and only if they are about disjointly realized curves, and sufficiently large positive powers of any finite collection of twists generate a right-angled Artin group by Theorem 1.1. Two curves in Σ are disjointly realized if and only if the corresponding vertices in $\mathcal{C}(\Sigma)$ are adjacent.

Conversely, let $A(\Gamma) \rightarrow \text{Mod}_{g,p}$ be an embedding as in the statement of the lemma. Identify the curve about which each vertex generator of $A(\Gamma)$ is a power of a twist, and identify the corresponding vertices in the curve complex. Filling in the edges between vertices which are adjacent in $\mathcal{C}(\Sigma)$ produces a copy of Γ in $\mathcal{C}(\Sigma)$. It follows that the obvious map $\Gamma \rightarrow \mathcal{C}(\Sigma)$ is a flag embedding. \square

8. EMBEDDINGS AND COHOMOLOGY

8.1. Commensurability questions. We first provide a proof that mapping class groups of genus $g \geq 3$ surfaces cannot be commensurable with right-angled Artin groups, using soft cohomological dimension arguments:

Proposition 8.1. *If $g > 2$ or if $g = 2$ and $p > 0$ then $\text{Mod}_{g,p}$ is not commensurable with a right-angled Artin group.*

Proof. Let Γ be a right-angled-Artin group such that $\text{Mod}_{g,p}$ is commensurable with Γ . Then $\text{Mod}_{g,p}$ and Γ have the same virtual cohomological dimension. It is a standard fact about right-angled Artin groups that their cohomological dimension is equal to the rank of a maximal abelian subgroup (equivalently this is the dimension of the Salvetti complex $\mathcal{S}(\Gamma)$, which is a finite-dimensional $K(\Gamma, 1)$). See [11], for instance. The work of Harer in [32] shows that if Σ is closed of genus g then $\text{vcd Mod}_g = 4g - 5$. On the other hand, the maximal abelian subgroup of Mod_g has rank $3g - 3$, so $\text{vcd } \Gamma = 3g - 3$. The only way $3g - 3 = 4g - 5$ is when $g = 2$.

If $p \neq 0$ then [32] shows that $\text{vcd Mod}_{g,p} = 4g + p - 4$. The rank of a maximal abelian subgroup of $\text{Mod}_{g,p}$ is $3g - 3 + p$. These two numbers are equal only when $g = 1$. \square

8.2. Virtual embedding questions. The question of whether a mapping class group virtually injects into right-angled Artin group is more subtle. Let $\text{Mod}^1(\Sigma)$ denote the mapping class group of Σ with one marked point. We have that $\text{Mod}^1(\Sigma)$ contains a copy of $\pi_1(\Sigma)$.

Lemma 8.2. *Let Σ be closed and let $H < \text{Mod}^1(\Sigma)$ be a finite index subgroup which contains $\pi_1(\Sigma)$. Then H is not residually torsion-free nilpotent. In particular, H does not inject into a right-angled Artin group.*

Proof. H virtually contains every group which is a semidirect product of the form

$$1 \rightarrow \pi_1(\Sigma) \rightarrow \Gamma \rightarrow \mathbb{Z} \rightarrow 1.$$

In particular, H virtually contains the fundamental group of each fibered hyperbolic manifold which admits a fibration with fiber Σ . It is possible to find elements of the symplectic group which act irreducibly on $H_1(\Sigma, \mathbb{Q})$, and all of whose powers act irreducibly on $H_1(\Sigma, \mathbb{Q})$. It follows that Γ^{ab} is cyclic. Since nilpotent groups with cyclic abelianization are themselves cyclic, we have that Γ is not residually torsion-free nilpotent. \square

Modifying the argument of Lemma 8.2 to show that mapping class groups are not virtually special seems difficult (see Propositions 8.4 and 8.5 below). To overcome these difficulties, we exploit certain non-virtually special subgroups of sufficiently complicated mapping class groups. Recall that a group G is called **residually finite rationally solvable** or **RFRS** if there exists an exhaustive filtration of normal (in G) finite index subgroups of G

$$G = G_0 > G_1 > G_2 > \dots$$

such that $G_{i+1} > \ker\{G_i \rightarrow G_i^{ab} \otimes \mathbb{Q}\}$. Agol proves that right-angled Artin groups are virtually RFRS (see [1]). Furthermore, any subgroup of a RFRS group is RFRS again, as follows easily from the definition.

Lemma 8.3. *Let Σ be a surface of genus $g \geq 2$ and one boundary component, and let $H < \text{Mod}(\Sigma)$ be a finite index subgroup. Then H is not RFRS. In particular, H does not inject into a right-angled Artin group.*

Proof. We have that $\text{Mod}(\Sigma)$ contains a copy of the fundamental group of the unit tangent bundle U of a genus two surface Σ_2 (see [23], though this fact is essentially due to J. Birman). We have that $\pi_1(U)$ fits into a non-split central extension

$$1 \rightarrow \mathbb{Z} \rightarrow \pi_1(U) \rightarrow \pi_1(\Sigma_2) \rightarrow 1.$$

Let Γ be the fundamental group of a nontrivial circle bundle of a closed surface Σ with trivial monodromy, and let $\{G_i\}$ be any candidate for an exhausting sequence which witnesses the RFRS condition on Γ . We will again show that $\{G_i\}$ cannot exhaust Γ . We view Γ as a central extension

$$1 \rightarrow \mathbb{Z} \rightarrow \Gamma \rightarrow \pi_1(\Sigma) \rightarrow 1.$$

Consider Γ^{ab} . Clearly there is a surjection $\Gamma^{ab} \rightarrow H_1(\Sigma, \mathbb{Z})$. Suppose that the central copy of \mathbb{Z} in Γ maps injectively to Γ^{ab} . We would then be able to describe Γ^{ab} as a central extension of the form

$$1 \rightarrow \mathbb{Z} \rightarrow \Gamma^{ab} \rightarrow H_1(\Sigma, \mathbb{Z}) \rightarrow 1,$$

and since Γ^{ab} is abelian the extension would have to split. In particular, its classifying cocycle in $H^2(\mathbb{Z}^{2g}, \mathbb{Z})$ would be trivial, where g is the genus of Σ . The pullback cocycle in $H^2(\Sigma, \mathbb{Z})$ would also be zero, so that the extension describing Γ itself would have to split, a contradiction. It follows that the central copy of \mathbb{Z} is in the kernel of the map $\Gamma \rightarrow \Gamma^{ab} \otimes \mathbb{Q}$. We therefore have $\mathbb{Z} < G_1$.

We inductively suppose that $\mathbb{Z} < G_i$. Write $H_i = G_i/\mathbb{Z}$, where we quotient by the central copy of \mathbb{Z} . Since H_i is a subgroup of $\pi_1(\Sigma)$, we get a restriction map $H^2(\pi_1(\Sigma), \mathbb{Z}) \rightarrow H^2(H_i, \mathbb{Z})$. Recall that since H_i has finite index in $\pi_1(\Sigma)$, we have that by general cohomology of groups there is a natural corestriction map $H^2(H_i, \mathbb{Z}) \rightarrow H^2(\pi_1(\Sigma), \mathbb{Z})$, and the composition of the corestriction with the restriction (often written $Cor \circ Res$) is multiplication by $[\pi_1(\Sigma) : H_i]$ (cf. [10]). It follows that if $e \neq 0$ then $Cor \circ Res(e) \neq 0$ so that $Res(e) \neq 0$. As in the base case, it follows that the central copy of \mathbb{Z} cannot map injectively into G_i^{ab} , so that this copy of \mathbb{Z} is in the kernel of the map $G_i \rightarrow G_i^{ab} \otimes \mathbb{Q}$. It follows that $\mathbb{Z} < G_{i+1}$. It follows that the original central copy of \mathbb{Z} was contained in each G_i , which is what we set out to prove. \square

Alternatively, if $\pi_1(U)$ was virtually RFRS then U would virtually fiber over the circle, which cannot happen. In some sense, the failure of the fundamental group of the unit tangent bundle of a surface to be RFRS is the algebraic phenomenon underlying the failure of the unit tangent bundle to fiber over the circle.

As a corollary to the considerations above, we can obtain Theorem 1.5:

Proof of Theorem 1.5. If Σ contains a genus two surface with one boundary component S as a subsurface then its mapping class group contains the mapping class group of S . In particular, $\text{Mod}(\Sigma)$ contains $\pi_1(U)$, the fundamental group of the unit tangent bundle of Σ_2 . But we have shown that such a mapping class group cannot virtually inject into a right-angled Artin group. \square

8.3. Virtual speciality questions. It is interesting to consider the cases of simpler mapping class groups for which our methods here do not work. If a given mapping class group is virtually special, then it must be virtually residually torsion-free nilpotent. Consider the following basic observations, which can be viewed as variations on Lemma 8.2:

Proposition 8.4. *Let $\phi \in \text{Aut}(F_n)$ be an automorphism and let*

$$1 \rightarrow F_n \rightarrow N_\phi \rightarrow \mathbb{Z} \rightarrow 1$$

be the semidirect product formed by taking the \mathbb{Z} -conjugation given by ϕ . If the action of ϕ has no fixed vector in $H_1(F_n^{ab}, \mathbb{Q})$, then N_ϕ is not residually torsion-free nilpotent.

Proof. The abelianization of N_ϕ factors through the semidirect product

$$1 \rightarrow F_n^{ab} \rightarrow \overline{N_\phi} \rightarrow \mathbb{Z} \rightarrow 1.$$

Since ϕ has no fixed vector in $H_1(F_n^{ab}, \mathbb{Q})$ then $N_\phi^{ab} \otimes \mathbb{Q}$ has rank one. It follows that any torsion-free nilpotent quotient of N_ϕ has cyclic abelianization. It is well-known that a nilpotent group with cyclic abelianization is itself cyclic, whence the claim. \square

Proposition 8.5. *Let $\Gamma < \text{Aut}(F_n)$ be a subgroup which contains $\text{Inn}(F_n)$ and let $\phi \in \Gamma$. Then $N_\phi < \Gamma$.*

Proof. Let $g \in F_n$ and let i_g denote conjugation by g . We have that $\phi \circ i_g \circ \phi^{-1} = i_{\phi(g)}$. Indeed,

$$\phi \circ i_g \circ \phi^{-1}(h) = \phi(g\phi^{-1}(h)g^{-1}) = \phi(g)h\phi(g)^{-1}.$$

It follows that the subgroup of Γ generated by ϕ and F_n is isomorphic to N_ϕ . \square

Both of these propositions hold with automorphisms of free groups replaced by mapping classes. Consider the set of hyperbolic 3-manifolds built by suspending pseudo-Anosov homeomorphisms in $\text{Mod}_{g,p}$ with fundamental groups $\{\Gamma_i\}$. Here, the meaning of $\text{Mod}_{g,p}$ is that we require the homeomorphisms and isotopies to preserve the punctures of $\Sigma_{g,p-1}$, together with an extra marked point. This way, we get an action of $\text{Mod}_{g,p}$ on $\pi_1(\Sigma_{g,p-1})$. In particular, whenever $\psi \in \text{Mod}_{g,p}$, we have that the fundamental group $\pi_1(M_\psi)$ of the mapping torus of ψ is contained in $\text{Mod}_{g,p}$.

If $\text{Mod}_{g,p}$ is virtually special then it follows that there is a universal finite index to which one could pass in any of the $\{\Gamma_i\}$ and obtain first Betti number greater than one. Indeed, suppose that $G < \text{Mod}_{g,p}$ embeds into a right-angled Artin group. Then $G \cap \Gamma_i$ is residually torsion-free nilpotent, so that $b_1(G \cap \Gamma_i) \geq 2$.

This observation seems strange at first, but consider why it holds for a once-punctured torus. Taking any nontrivial abelian cover, there is a finite index subgroup of the mapping class group which preserves the homology class of each lift of the puncture, so that there is a lot of invariant homology in any suspension of any mapping class.

We close with the following conjectures:

Conjecture 8.6. *The following groups are virtually special:*

- (1) *The mapping class group Mod_2 .*
- (2) *The braid groups B_n , for all n .*

REFERENCES

- [1] Ian Agol. Criteria for virtual fibering. *J. Topol.*, 1, 269–284, 2008.
- [2] G. Baumslag and J.E. Roseblade. Subgroups of direct products of free groups, *J. London Math. Soc.* 30, 44–52, 1984.
- [3] Jason Behrstock and Ruth Charney. Divergence and quasi-morphisms of right-angled Artin groups. Preprint.
- [4] Ricardo Benedetti and Carlo Petronio. *Lectures on hyperbolic geometry*. Universitext, Springer, 1992.
- [5] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.* 129, 123–139, 1997.

- [6] Joan S. Birman and Alex Lubotzky and John McCarthy. Abelian and solvable subgroups of the mapping class groups. *Duke Math. J.* 50, no. 4, 1107–1120, 1983.
- [7] Martin R. Bridson. Semisimple actions of mapping class groups on CAT(0) spaces. To appear in *The Geometry of Riemann Surfaces*, LMS Lecture Notes 368.
- [8] Martin R. Bridson. On the dimension of CAT(0) spaces where mapping class groups act. Preprint, 2009.
- [9] Martin R. Bridson and Charles F. Miller III. Structure and finiteness properties of subdirect products of groups. To appear in *Proc. London Math. Soc.*
- [10] Kenneth S. Brown. *Cohomology of groups*. Graduate Texts in Mathematics, no. 87, Springer, New York, 1982.
- [11] Ruth Charney. An introduction to right-angled Artin groups. *Geom. Dedicata*, 125, 141–158, 2007.
- [12] Ruth Charney and Michael Farber. Random groups arising as graph products. Preprint.
- [13] Matt Clay, Chris Leininger and Johanna Mangahas. The geometry of right angled Artin subgroups of mapping class groups. To appear in *Groups, Geom. Dyn.*
- [14] John Crisp and Benson Farb. The prevalence of surface groups in mapping class groups. Preprint.
- [15] John Crisp and Luis 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, no. 1, 19–36, 2001.
- [16] John Crisp and Bert Wiest. Quasi-isometrically embedded subgroups of braid and diffeomorphism groups. *Trans. Amer. Math. Soc.* 359, no. 11, 5485–5503, 2007.
- [17] Matthew B. Day. Peak reduction and finite presentations for automorphism groups of right-angled Artin groups. *Geom. Topol.* 13, 817–855, 2009.
- [18] Pierre de la Harpe. *Topics in Geometric Group Theory*. University of Chicago Press, 2000.
- [19] Carl Droms. Subgroups of graph groups. *J. Algebra* 110, 519–522, 1987.
- [20] Carl Droms. Isomorphisms of graph groups. *Proc. Amer. Math. Soc.* 100, no. 3, 407–408, 1987.
- [21] Benson Farb. Some problems on mapping class groups and moduli space. In *Problems on Mapping Class Groups and Related Topics. Proc. Symp. Pure and Applied Math.*, Volume 74, 2006.
- [22] Benson Farb and Alexander Lubotzky and Yair Minsky. Rank one phenomena for mapping class groups. *Duke Math. J.*, Vol. 106, No. 3, 581–597, 2001.
- [23] Benson Farb and Dan Margalit. *A Primer on the mapping class group*. Princeton Mathematical Series, 49. Princeton University Press, Princeton, NJ, 2012.
- [24] A. Fathi and F. Laudenbach and V. Poénaru. *Travaux de Thurston sur les Surfaces*. Soc. Math. de France, Paris, Astérisque 66-67, 1979.
- [25] Edward Formanek and Claudio Procesi. The automorphism group of a free group is not linear. *J. Algebra* 149, no. 2, 494–499, 1992.
- [26] Koji Fujiwara. Subgroups generated by two pseudo-Anosov elements in a mapping class group. I. Uniform exponential growth. *Groups of diffeomorphisms*, 283–296, *Adv. Stud. Pure Math.*, 52, Math. Soc. Japan, Tokyo, 2008.
- [27] Louis Funar. On power subgroups of mapping class groups. Preprint.
- [28] William M. Goldman. *Complex hyperbolic geometry*. Oxford Mathematical Monographs, 1999.
- [29] Joseph Gubeladze. The isomorphism problem for commutative monoid rings. *J. Pure Appl. Algebra*, 129(1):35–65, 1998.
- [30] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *GAFSA, Geom. funct. anal.* 17, 1551–1620, 2008.
- [31] Ursula Hamenstädt. Geometric properties of the mapping class group. In *Problems on Mapping Class Groups and Related Topics. Proc. Symp. Pure and Applied Math.*, Volume 74, 2006.
- [32] John L. Harer. The virtual cohomological dimension of the mapping class group of an orientable surface. *Invent. Math.* 84, no. 1, 157–176, 1986.
- [33] Nikolai Ivanov. Subgroups of Teichmüller modular groups. *Translations of Mathematical Monographs*, 115, 1992.
- [34] Tadeusz Januszkiewicz and Jacek Świątkowski. Hyperbolic Coxeter groups of large dimension. *Comment. Math. Helv.* 78, 555–583, 2003.
- [35] Michael Kapovich and Bernhard Leeb. Actions of discrete groups on nonpositively curved spaces. *Math. Ann.* 306, no. 2, 341–352, 1996.
- [36] Svetlana Katok. *Fuchsian groups*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1992.
- [37] Sang-hyun Kim and Thomas Koberda. Embedability between right-angled Artin groups. Preprint 2011.
- [38] Thomas Koberda. Ping-pong lemmas with applications to geometry and topology. To appear in *IMS Lecture Notes*, Singapore, 2012.
- [39] Michael R. Laurence. A generating set for the automorphism group of a graph group. *J. London Math. Soc.* (2) 52 no. 2, 318–334, 1995.
- [40] Gilbert Levitt. Unsolvability of the isomorphism problem for [free abelian]-by-free groups. Preprint.
- [41] Michael Lönne. Presentations of subgroups of the braid group generated by powers of band generators. *Topology Appl.* 157, no. 7, 1127–1135, 2010.

- [42] Chris Leininger and Alan Reid. A combination theorem for Veech subgroups of the Mapping class group. *Geom. and Funct. Anal.* 16, 403–436, 2006.
- [43] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer, New York, 1977.
- [44] K.A. Mihailova. The occurrence problem for free products of groups. (Russian) *Mat. Sb. (N.S.)* 75 (117) 199–210, 1968.
- [45] John McCarthy and Athanase Papadopoulos. Dynamics on Thurston’s sphere of projective measured foliations. *Comment. Math. Helv.* 64, no. 1, 133–166, 1989.
- [46] Charles F. Miller, III. *On group-theoretic decision problems and their classification*. *Ann. of Math. Studies*, 68, 1971.
- [47] Charles F. Miller III. Decision problems for groupssurvey and reflections, in *Algorithms and classification in combinatorial group theory* (Berkeley, CA, 1989), *Math. Sci. Res. Inst. Publ.*, 23, Springer–Verlag, 1–59, 1992.
- [48] Hee Oh. Discrete subgroups generated by lattices in opposite horospherical subgroups. *J. Algebra* 203, 621–676, 1998.
- [49] R.C. Penner and J.L. Harer. *Combinatorics of train tracks*. *Annals of Mathematics Studies*, no. 125, *Princeton University Press*, 1992.
- [50] Lucas Sabalka. Embedding right-angled Artin groups into graph braid groups. *Geometriae Dedicata* 124, no. 1, 191–198, 2007.
- [51] Lucas Sabalka. On rigidity and the isomorphism problem for tree braid groups. *Groups, Geometry, and Dynamics*, 3(3):469–523, 2009.
- [52] H. Servatius. Automorphisms of graph groups. *J. Algebra* 126, no. 1, 34–60, 1989.
- [53] John Stallings. The topology of finite graphs. *Invent. Math.* 71, no. 3, 551–565, 1983.
- [54] John Stillwell. The occurrence problem for mapping class groups. *Proc. Amer. Math. Soc.* 101, no. 3, 411–416, 1987.
- [55] Jacques Tits. Free subgroups in linear groups. *J. Algebra*, 20, 250–270, 1972.
- [56] Stephen Wang. Representations of surface groups and right-angled Artin groups in higher rank. *Algebr. Geom. Topol.* 7, 1099–1117, 2007.

DEPARTMENT OF MATHEMATICS, HARVARD UNIVERSITY, 1 OXFORD ST., CAMBRIDGE, MA 02138

E-mail address: koberda@math.harvard.edu