RIGHT-ANGLED ARTIN GROUPS AND THEIR SUBGROUPS

THOMAS KOBERDA

Abstract. These are notes for a course offered at Yale University in the spring semester of 2013.

March 7, 2013

Contents

1991 *Mathematics Subject Classification.* Primary ; Secondary .

1. Introduction

In this course, we will be discussing right-angled Artin groups and their subgroups. Perhaps a more complete title for the course would be, "Right-angled Artin groups and their subgroups from an algebraic and geometric point of view, and their relationship with mapping class groups".

1.1. **Right-angled Artin groups.**

In the sequel we will use Γ to denote a finite *simplicial graph.* We will write $V = V(\Gamma)$ for the finite set of vertices and $E = E(\Gamma) \subset V \times V$ for the set of edges, viewed as unordered pairs of vertices. The requirement that Γ be simplicial simply means that the diagonal of $V \times V$ is excluded from the set of edges. The *right-angled Artin group* on Γ is the group

$$A(\Gamma) \cong \langle V \mid [v_i, v_j] = 1 \text{ whenever } (v_i, v_j) \in E \rangle.$$

In other words, $A(\Gamma)$ is generated by the vertices of Γ, and the only relations are given by commutation of adjacent vertices.

1.2. **The universal property of right-angled Artin groups.**

Algebraically, a right-angled Artin group is the universal group with specified commutation and non–commutation among its vertices. For any subset $S \subset G$ of a group, we build the *commutation graph* of S, written $Comm(S)$, as follows. The vertices of $Comm(S)$ are the elements of S, and two vertices of S are connected by an edge if they commute in G. The following proposition gives the universal property of right-angled Artin groups.

Proposition 1.1. *Let G be a group and let $S \subset G$ be a finite subset. The inclusion $S \to G$ extends to a unique homomorphism*

$$A(Comm(S)) \to G$$

which agrees with the identification $V(Comm(S)) \cong S$.

In the universal property, we required S to be finite because right-angled Artin groups are defined to be finitely generated. One could extend the definition to accommodate right-angled Artin groups on arbitrary graphs, whereupon the finiteness hypothesis on S would be unnecessary.

1.3. **Other Artin groups and Coxeter groups.**

We will briefly digress to mention the relationship between right-angled Artin groups, other Artin groups, and Coxeter groups. In a group G, two elements s and t satisfy an *Artin relation of length n* if

$$sts \cdots = tst \cdots ,$$

where on both sides there are exactly n letters occurring, alternating between s and t. Observe that an Artin relation of length two is just commutation. If s and t satisfy no relations, we say that they satisfy an *Artin relation of length infinity.*

Let Γ be a finite *labelled simplicial graph*, which is two say a finite simplicial graph where each edge is labelled by an integer $2 \leq n < \infty$. The *Artin group* on Γ, writted $A(\Gamma)$, is the group generated by the vertices of Γ, and two vertices v and w connected by an edge labelled n satisfy an Artin relation of length n. The vertices of Γ are a distinguished set of generators and are often called an *Artin system*. We will adopt the convention that vertices not connected by an edge in Γ are connected by an invisible edge labelled by infinity, so that a right-angled Artin group is an Artin group where all the edge labels are two.

Occasionally, the convention is adopted that vertices in Γ which are not connected by an edge are connected by an invisible edge whose label is 2. Thus, is P_n is a path on n vertices with all the edges labels equal to three, the resulting Artin group would be

$$B_n = \langle s_1, \ldots, s_n \mid s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1},\ [s_i, s_j] = 1 \text{ whenever } |i-j| \geq 2 \rangle,$$

which is the usual planar braid group.

Every Artin group has an associated *Coxeter group*, written $C(\Gamma)$. If V is an Artin system for an Artin group, the Coxeter group is obtained from the Artin group by imposing the relations $v^2 = 1$ for each $v \in V$. In the context of Coxeter groups, the set V is often called a *Coxeter system* for the Coxeter group. Observe that the relation $v^2 = 1$ for each $v \in V$ allows us to rewrite an Artin relation of length n between v and w as $(vw)^{2n} = 1$. A general Coxeter group on a labelled graph Γ is generated by V together with the relation $(vw)^n = 1$ whenever v and w are connected by an edge labelled by n.

Coxeter groups are intimately related with reflection groups. The basic observation is that if two hyperplanes H_1 and H_2 meet at an angle π/k and if r_1 and r_2 denote reflections across H_1 and H_2 respectively, r_1 ad r_2 satisfy the relations $r_1^2 = r_2^2 = 1$ and $(r_1 r_2)^k = 1$. The lattermost of these relations results from the fact that $r_1 r_2$ is a rotation by $2\pi/k$. It is true (though not trivial) that if $\{H_1, \ldots, H_n\}$ is a hyperplane configuration with each pair (H_i, H_j) of distinct hyperplanes meeting at an angle of the form $\pi/k_{i,j}$ with $k_{i,j} \geq 3$, then the group generated by reflections in these hyperplanes is isomorphic to the Coxeter group on a graph Γ whose vertices are labelled by $\{H_1, \ldots, H_n\}$ and whose edges are given by nonempty intersection and are labelled by $k_{i,j}$.

Coxeter groups give an explanation for the choice of terminology "right-angled Artin group". Let $\{H_1, \ldots, H_n\}$ be a hyperplane configuration where for each $i \neq j$ we have H_i and H_j are disjoint or meet at a right angle. Viewing $\{H_1, \ldots, H_n\}$ as an abstract Coxeter system, we have that two elements v and w of the Coxeter system either satisfy no relations or satisfy $(vw)^2 = 1$. This latter relation can be rewritten as $vw = wv$. It follows that the resulting Coxeter group is the Coxeter group associated to the right-angled Artin group on $\{H_1, \ldots, H_n\}$, where edges are given by nonempty intersection.

1.4. **Goals of the course.** We will now briefly outline the topics we will cover in this course:

(1) Basic theory of right-angled Artin groups. Ping–pong lemmas, two–generator subgroups, structures of centralizers, Salvetti complexes, automorphisms.
(2) "Classical" subgroups of right-angled Artin groups. Surface subgroups, one–ended subgroups, Bestvina–Brady subgroups, general combinatorial Morse theory.
(3) Right-angled Artin groups in mapping class groups. General theory of embedding right-angled Artin groups in mapping class groups and the "up to powers" classification of finitely generated subgroups of mapping class groups following Koberda and Kim–Koberda, curve complex methods and quasi–isometrically embedded right-angled Artin subgroups of mapping class groups following Clay–Leininger–Mangahas.
(4) The unsolvability of the isomorphism problem for finitely presented subgroups of right-angled Artin groups and mapping class groups.
(5) Applications of mapping class groups to the study of right-angled Artin groups. Right-angled Artin subgroups of right-angled Artin groups following Kim–Koberda.

(6) Curve complex theory for right-angled Artin groups. The Bowditch and Masur–Minsky machinery for right-angled Artin groups following Kim–Koberda.

The reader may find that there are some glaring omissions from our discussion of subgroups of right-angled Artin groups. Most notably, we will avoid a discussion of special cube complexes (following Haglund, Wise and others) and of hyperbolic 3–manifold subgroups of right-angled Artin groups (following Wise and Agol–Groves–Manning). The reasons for this omission are several–fold, though mostly we avoid discussing these topics for the sake of brevity, because many of the methods and the flavors of the results are somewhat different from the topics at hand, and because the author feels he could not do the subject full justice.

2. Basic tools

We will first describe some of the basic tools for describing right-angled Artin groups and their subgroups.

2.1. Concrete examples of right-angled Artin groups.

From a certain perspective, right-angled Artin groups interpolate between free groups and free abelian group. Let us consider a few examples:

(1) $\Gamma = K_n$, the complete graph on n vertices. Then $A(\Gamma) \cong \mathbb{Z}^n$.
(2) $\Gamma = D_n$, the completely disconnected graph on n vertices. Then $A(\Gamma) \cong F_n$.
(3) $\Gamma = P_3$, the path on three vertices. Then $A(\Gamma) \cong F_2 \times \mathbb{Z}$.
(4) $\Gamma = C_4$, the cycle of length four (a square). Then $A(\Gamma) \cong F_2 \times F_2$.

For more general graphs, there are usually no other more familiar names one can give right-angled Artin groups – the description $A(\Gamma)$ is the most efficient one. It is generally true that graph theoretic properties of Γ translate into algebraic properties of $A(\Gamma)$. Some correspondences are obvious, whereas others are less so:

(1) The group $A(\Gamma)$ splits as a nontrivial free product if and only if Γ is disconnected.
(2) The group $A(\Gamma)$ splits as a nontrivial direct product if and only if Γ splits as a nontrivial join. Recall that a graph is a nontrivial join if its complement is disconnected.
(3) Let P_4 denote a path on four vertices. The group $A(\Gamma)$ contains a copy of $A(P_4)$ if and only if Γ contains P_4 as an induced subgraph.
(4) The group $A(\Gamma)$ contains a copy of $F_2 \times F_2$ if and only if Γ contains an induced square.

As the course progresses, we will give proofs of items (3) and (4) above. The following is an interesting and generally open problem, which is in the spirit of the four items above:

Question 2.1. *Under what conditions does $A(\Gamma)$ contain a closed hyperbolic surface group?*

This question will be a major subject of discussion in a later part of the course.

2.2. Topological models for right-angled Artin groups.

Right-angled Artin groups have very easy to describe $K(G,1)$ spaces, which allow for a straightforward calculation of the cohomology ring $H^*(A(\Gamma), \mathbb{Z})$.

For each $n > 0$ we will write T^n for an n–torus given by $\mathbb{R}^n/\mathbb{Z}^n$. A fundamental domain for T^n is given by the unit cube in the octant where each coordinate is positive. The sub–cubes given by setting certain coordinates to zero give *distinguished subtori* of T^n. For each k, we will fix the "canonical" isomorphism of T^k with any subtorus of T^n given by setting $n-k$ of the coordinates to be zero. Specifically, we will think of the k coordinates for T^k as n coordinates, and merely insert zeros to coincide with the $n-k$ zeros of the subtorus of T^n. The subtorus given by varying k coordinates and setting the other $n-k$ to zero is the k–torus *spanned* by those k coordinates.

We build the Salvetti complex $S(\Gamma)$ associated to Γ by adding higher and higher dimensional tori. We start with a single vertex x_0 which will be the zero skeleton $S^{(0)}(\Gamma)$. For each vertex of Γ, we attach a copy of T^1, identifying x_0 with the distinguished subtorus of T^1. Inductively, suppose we have constructed the n–skeleton $S^{(n)}(\Gamma)$. To build the $(n+1)$–skeleton, we take one copy of T^{n+1} for each complete subgraph $K \subset \Gamma$ on $n+1$ vertices in Γ, and we label the coordinate directions in T^{n+1} with the vertices in K. We glue T^{n+1} to $S^{(n)}(\Gamma)$ by identifying the torus corresponding to a subgraph of K with the subtorus of T^{n+1} spanned by the corresponding coordinates.

Proposition 2.2. *The complex $S(\Gamma)$ has the structure of a finite CW–complex and satisfies*

$$\pi_1(S(\Gamma)) \cong A(\Gamma).$$

The universal cover of $S(\Gamma)$ is contractible, so that $S(\Gamma)$ is a $K(A(\Gamma), 1)$.

Proof. Exercise. □

The preceding description of the Salvetti complex gives $S(\Gamma)$ the structure of a locally CAT(0) cube complex, so that the right-angled Artin group $A(\Gamma)$ acts properly, discontinuously, cocompactly, and by isometries on a CAT(0) cube complex. Presently, we will not dwell on this aspect of right-angled Artin group theory.

The explicit nature of $S(\Gamma)$ and its role as $K(A(\Gamma), 1)$ makes it fairly straightforward to compute $H^*(A(\Gamma), \mathbb{Z})$. Evidently, $H^1(A(\Gamma), \mathbb{Z})$ is spanned by classes $V^* = \{v_1^*, \ldots, v_k^*\}$ which are dual to the 1–tori in $S^{(1)}(\Gamma)$, which in turn are in bijective correspondence with $V = V(\Gamma)$. If $V_0 \subset V$ span a complete subgraph of Γ then the corresponding dual classes V_0^* span an integral exterior algebra $\Lambda_{\mathbb{Z}}(V_0^*) \subset H^*(A(\Gamma), \mathbb{Z})$. Furthermore, if $(v, w) \notin E$, we have that $v^* \cup w^* = 0$. These computations are sufficient to compute $H^*(A(\Gamma), \mathbb{Z})$:

Proposition 2.3. *Let Γ be a finite simplicial graph. Then $H^*(A(\Gamma), \mathbb{Z})$ is generated by the degree one elements V^*, and*

$$H^*(A(\Gamma), \mathbb{Z}) \cong \left(\bigoplus_{V_0^* \subset V^*} \Lambda_{\mathbb{Z}}(V_0^*) \right) / \sim,$$

where V_0^* ranges over subsets of V^* which are dual to complete subgraphs of Γ, and where the generators of two exterior algebras are identified if they are equal in V^*.

Proof. Exercise. □

Proposition 2.3 has some useful though elementary corollaries:

Corollary 2.4. *The cohomological dimension of $A(\Gamma)$ is equal to the size of the largest complete subgraph of Γ.*

Corollary 2.5. *The group $A(\Gamma)$ is not isomorphic to the fundamental group of a closed, orientable manifold unless Γ is complete.*

Corollary 2.6. *The group $A(\Gamma)$ is torsion–free.*

2.3. **Graph isomorphism versus group isomorphism.** The next fact we would like to establish is the bijective correspondence between isomorphism types of right-angled Artin groups and isomorphism types of finite simplicial graphs. In general, it appears to be unknown when two labelled simplicial graphs give rise to isomorphic Artin groups, but the right-angled case is significantly easier.

Theorem 2.7. *Two right-angled Artin groups $A(\Gamma)$ and $A(\Lambda)$ are isomorphic if and only if $\Gamma \cong \Lambda$.*

This result was originally proven by Droms in [18], and relied on an older result which asserts that two graph algebras over a field are isomorphic if and only if the underlying graphs are isomorphic (see [28]). The *graph algebra* over a field K is written $K(\Gamma)$ and is the free algebra on the vertices of Γ, subject to the relations $[v_i, v_j] = 0$ whenever $(v_i, v_j) \in E$. Droms shows that if $A(\Gamma) \cong A(\Lambda)$ then $K(\Gamma) \cong K(\Lambda)$, thus establishing the result.

Theorem 2.7 has been recovered by various authors, including Sabalka in [40] and Koberda in [32]. We will give an argument following the lines of [32].

Proof of Theorem 2.7. We will give the proof in the case where Γ has no triangles and no squares, with the details of the general case being left to the reader. If $\Gamma \cong \Lambda$ then $A(\Gamma) \cong A(\Lambda)$. Conversely, suppose that $A(\Gamma) \cong A(\Lambda)$. Then

$$H^*(A(\Gamma), \mathbb{Q}) \cong H^*(A(\Lambda), \mathbb{Q}).$$

Write $W = H^1(A(\Gamma), \mathbb{Q})$. If $w \in W$ is a cohomology class, we will consider the assignment

$$w \mapsto f_w,$$

where

$$f_w : H^1(A(\Gamma), \mathbb{Q}) \to H^2(A(\Gamma), \mathbb{Q})$$

is defined by $f_w(x) = w \cup x$.

In general, w can be viewed as a rational linear combination

$$\sum_{i=1}^{n} a_i v_i^*$$

of cohomology classes dual to the vertices of Γ, and if

$$x = \sum_{i=1}^{n} b_i v_i^*$$

is another such combination, we have

$$f_w(x) = \sum_{i=1}^{n} \sum_{j=1}^{n} (a_i b_j)(v_i^* \cup v_j^*),$$

where $v_i^* \cup v_j^*$ can be viewed as dual to the edge connecting v_i and v_j if it exists, and zero otherwise. For a cohomology class

$$w = \sum_{i=1}^{n} a_i v_i^*,$$

we will call the vertices dual to the basis elements for which the coefficients are nonzero the *support* of w. Let $W_0 \subset W$ be the subspace characterized by $w \in W_0$ if and only if the rank of f_w is zero. Evidently w is supported on degree zero vertices, so that Γ has $\dim W_0$ isolated vertices. We will quotient W by W_0 and call the result W again, by abuse of notation.

Now consider $w \in W$ such that the rank of f_w is exactly one. Observe that w must be supported on degree one vertices, and conversely any cohomology class w supported on exactly one degree one vertex has f_w of rank one. Thus, the span of all such cohomology classes $W_1 \subset W$ has dimension equal to the number of degree one vertices in Γ.

Suppose that f_w has rank exactly $n \geq 1$. Observe that w cannot be supported on a vertex of degree more than n, by rank considerations. Conversely, every vertex v of degree n has f_{v^*} of rank n. By induction, we have that the dimension of the subspace W_n of W spanned by w such that f_w has rank at most n modulo the subspace W_{n-1} spanned by w such that f_w has rank at most $n-1$ is equal to the number of degree n vertices of Γ.

Suppose that f_w has rank exactly n, and that the support of w contains two vertices of degree n, say v_1 and v_2. We have that $\mathrm{St}(v_1) \cup \mathrm{St}(v_2)$ cannot contain more than n distinct vertices. Since v_1 and v_2 both have degree n, the only way that this can happen is if the links of v_1 and v_2 agree, or if v_1 and v_2 are adjacent and if

$$\mathrm{St}(v_1) \setminus \mathrm{St}(v_2) = \mathrm{St}(v_2) \setminus \mathrm{St}(v_1).$$

If the degrees of v_1 and v_2 are both at least two, then both of these possibilities are ruled out by the absence of squares and triangles in Γ. A similar argument shows that if v_1 has degree n and v_2 has degree between 2 and $n-1$, then w cannot be supported on both v_1 and v_2.

Thus, if $w \in W_n \setminus W_{n-1}$ with $n \geq 2$, we have that w is supported on a unique vertex of degree n. It follows that vertices of degree at least two can be identified uniquely as certain one–dimensional subspaces of W, possibly up to an element of W_1, and adjacency can be tested through the cup product pairing. Adjacency to degree one vertices can easily be determined after all adjacencies of vertices of degree at least two is determined. □

2.4. **Ping–pong and its applications.** The classical ping–pong lemma gives a method to verify that a particular action of a free group is faithful. There is an analogous ping–pong lemma for right-angled Artin groups which is also useful for analyzing both right-angled Artin groups and their subgroups.

Classical ping–pong was formulated by Klein in [29]:

Theorem 2.8. *Let $F_n = \langle x_1, \ldots, x_n \rangle$ be a free group of rank n which acts on a set Y. Let $\{Y, \ldots, Y_n\}$ be subsets of Y and let*

$$y_0 \in Y \setminus \bigcup_{i=1}^{n} Y_i.$$

Suppose that the following two conditions hold:

(1) *For all $k \neq 0$ and all i, we have $x_i^k(y_0) \in Y_i$.*
(2) *For all $i \neq j$ and all $k \neq 0$, we have $x_i^k(Y_j) \subset Y_i$.*

Then the action of F_n on Y is faithful.

Proof. Let $1 \neq w \in F_n$ be a reduced word. Then without loss of generality, we may assume that w can be written in the form $w = x_1^k \cdot w'$, where $k \neq 0$, where w' is a reduced word of length $\ell(w) - k$, and where w' cannot be written as $x_1^{\pm 1} w''$ with $\ell(w') > \ell(w'')$. It follows that $w(y_0) \in Y_1$ whereas $y_0 \notin Y_1$, so that the action of w is nontrivial. □

Observe that ping–pong does not require the sets $\{Y_1, \ldots, Y_n\}$ to be disjoint, only that their union is not all of Y. One can use ping–pong to prove various geometric statements such as the following:

Proposition 2.9. *For all $n \geq 2$, the matrices*

$$\begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ n & 1 \end{pmatrix}$$

generate a free subgroup of $SL_2(\mathbb{Z})$.

Proof. Exercise. □

We will now adapt the ping–pong lemma for right-angled Artin groups. The statement of the following ping–pong lemma can be found originally in a preprint of Crisp and Farb in [13] (see also [32] and [33]):

Theorem 2.10. *Let $A(\Gamma)$ and let X be a set with a $A(\Gamma)$–action. Suppose the following hold:*

(1) *There exist subsets $X_i \subset X$ for each vertex v_i of Γ whose union is properly contained in X.*
(2) *For each nonzero $k \in \mathbb{Z}$ and $(v_i, v_j) \notin E$, $v_i^k(X_j) \subset X_i$.*
(3) *For each nonzero $k \in \mathbb{Z}$ and $(v_i, v_j) \in E$, $v_i^k(X_j) \subset X_j$.*
(4) *There exists*

$$x_0 \in X \setminus \bigcup_{i \in V} X_i$$

such that for each nonzero $k \in \mathbb{Z}$, $v_i^k(x_0) \in X_i$.

Then the $A(\Gamma)$–action on X is faithful.

First proof of Theorem 2.10. As usual, we show that the action is faithful by induction on the length of words in $A(\Gamma)$, using V as a generating set. The base case is resolved by the conditions on x_0. For the inductive step, write $w = v_i^k w'$, where k is maximal in absolute value and w' has length strictly less than the length of w. By the inductive hypothesis, $w'(x_0) \in X_j$ for some j. If (v_i, v_j) does not form an edge of Γ then $g_i^k(w'(x_0)) \in X_i$. Otherwise, (v_i, v_j) forms an edge of Γ, so that $v_i^k(w'(x_0)) \in X_j$. In particular, $w(x_0) \neq x_0$. □

Note that when Γ is a completely disconnected graph then Theorem 2.10 restricts to the usual ping–pong lemma. It is interesting to note what Theorem 2.10 says in the case when Γ is a complete graph, so that $A(\Gamma)$ is a free abelian group. It says that there are subsets of X, one for each vertex, which are disjoint from the basepoint and which are each preserved by the vertices of Γ. When we consider n linearly independent Euclidean translations acting on $\mathbb{R}^n$, such subsets can be taken to be certain neighborhoods of infinity which do not contain the origin, for instance.

We now give a proof of Theorem 2.10 which was pointed out by M. Kapovich. The argument gives a slightly different perspective on right-angled Artin groups and ping–pong:

Second proof of Theorem 2.10. 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_0 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 $w(x_0)$ 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 3 and 4 of the assumptions on $A(\Gamma)$ and X gives the claim. Now we consider $w_n(w_{n-1} \cdots w_1(x_0))$. By induction, $w_{n-1} \cdots w_1(x_0)$ 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 conditions 2 and 3 imply that $w_n(X_i) \subset X_j$, so that $w(x_0) \in X_j$. In particular, $w(x_0) \neq x_0$. □

Note the following easy corollaries of Theorem 2.10:

Corollary 2.11. *Let $v_1, \ldots, v_k$ be vertices of the graph Γ, viewed as generators of $A(\Gamma)$, and let $n_1, \ldots, n_k$ be nonzero integers. Then*

$$\langle v_1^{n_1}, \ldots, v_k^{n_k} \rangle \cong A(\Gamma).$$

Let $V_0 \subset V$ be a subset of the vertices of Γ and let $\Gamma_0 \subset \Gamma$ be the subgraph spanned by V_0. The subgroup of $A(\Gamma)$ spanned by V_0 is called a *standard subgroup* of $A(\Gamma)$.

Corollary 2.12. *The standard subgroup generated by V_0 is isomorphic to $A(\Gamma_0)$.*

The following is a somewhat less trivial corollary of Theorem 2.10:

Proposition 2.13. *Suppose $A(\Gamma)$ acts as a faithful, discrete group of isometries of hyperbolic n–space $\mathbb{H}^n$. Then $A(\Gamma)$ splits as a free product of free abelian groups of rank at most $n - 1$.*

Proof. Exercise. □

2.5. **Centralizers.** We would like to understand the structure of centralizers in right-angled Artin groups. A reduced word $w \in A(\Gamma)$ is called *cyclically reduced* if it has minimal length among all reduced words given by cyclically permuting the letters of w. A subgroup of $A(\Gamma)$ is called a *join subgroup* if it is a standard subgroup generated of $A(\Gamma)$ by a subgraph of Γ which splits as a nontrivial join. The following result appears in [3], and we reproduce it and its proof here. We will not give the proof of Servatius' result, and instead we will refer to his paper.

Proposition 2.14 (Centralizer Theorem)**.** *Let $1 \neq w \in A(\Gamma)$ be a cyclically reduced word. The following are equivalent:*

1. *The element $w \in A(\Gamma)$ is contained in a join subgroup.*
2. *The centralizer of w in $A(\Gamma)$ is not cyclic.*
3. *The centralizer of w in $A(\Gamma)$ is contained in join subgroup.*

Proof. (1) implies (2) since the centralizer of $(g_1, g_2) \in G_1 \times G_2$ contains $C(g_1) \times C(g_2)$, which contains $\mathbb{Z}^2$. (3) implies (1) is trivial.

Let $S \subset V$ be any subset and let $lk(S)$ denote the set of vertices which are a distance 1 from each vertex of S. Observe that $lk(S)$ may be empty. In [41], Servatius proves that the centralizer of a cyclically reduced word w lies in the subgroup generated by the support $supp(w)$ of w and by $lk(supp(w))$, and that the centralizer is cyclic unless $lk(supp(w))$ is nonempty or if $supp(w)$ decomposes as a join. We have that if the centralizer of w is not cyclic then $supp(w) \cup lk(supp(g))$ spans a join in Γ, so that (2) implies (3). □

The following corollary is immediate.

Corollary 2.15. *If $A(\Gamma)$ splits as a nontrivial direct product then Γ splits as a nontrivial join.*

2.6. **Automorphisms.** In this subsection, we will note Laurence's Theorem ([34]), which is a resolution of Servatius' conjecture in [41] on the automorphisms of right-angled Artin groups. Servatius noted the following types of automorphisms of right-angled Artin groups and conjectured that they generated all of $\mathrm{Aut}(A(\Gamma))$:

1. Graph symmetries.
2. Inversions of vertex generators.
3. Partial conjugations. Let $v \in V$, and suppose the star $\mathrm{St}(v)$ of v separates Γ. Let Γ_0 be a union of connected components of $\Gamma \setminus \mathrm{St}(v)$, and let $u = v^{\pm 1}$. A partial conjugation by u is given by sending $x \in V(\Gamma)$ to uxu^{-1} whenever $x \in \Gamma_0$ and to x otherwise.
4. Dominated transvections. A vertex w dominates a vertex v whenever the link of v is contained in the star of w. Observe that a vertex always dominates itself. For distinct vertices v and w, the dominated transvection τ_w sends v to vw or wv and preserves the other vertices of Γ.

It is easy to check that the four types of maps defined on the vertices of Γ define automorphisms of $A(\Gamma)$. These automorphisms of $A(\Gamma)$ are called *elementary automorphisms* of $A(\Gamma)$.

Theorem 2.16 (Laurence)**.** *The elementary automorphisms of $A(\Gamma)$ generate $\mathrm{Aut}(A(\Gamma))$.*

Corollary 2.17. *If Γ has no separating stars and no domination relations between pairs of vertices then $\mathrm{Out}(A(\Gamma))$ is finite.*

Not every finite order outer automorphism comes from a graph symmetry composed with an inversion. Let Γ be the graph obtained from doubling a pentagon C_5 along one edge. Choose a vertex v on the doubled edge. Since the star of v separates Γ, we can define the partial conjugation by v, which conjugates only one component of $\Gamma \setminus \mathrm{St}(v)$. Let ϕ be the automorphism given by composing this partial conjugation with a reflection about the doubled edge. Observe that ϕ acts nontrivially on $H_1(A(\Gamma), \mathbb{Z})$, but that ϕ^2 is conjugation by v.

Proposition 2.18. *The automorphism ϕ above has finite order in $\mathrm{Out}(A(\Gamma))$ and does not coincide with a graph symmetry composed with a vertex inversion.*

The absence of domination has the following curious consequence, which we will appeal to again in the sequel:

Proposition 2.19. *Let Γ be a finite simplicial graph with no triangles, no squares and no degree one vertices. Suppose $1 \neq w \in A(\Gamma)$ is a cyclically reduced word with nonabelian centralizer. Then there is a nonzero k and a $v \in V(\Gamma)$ such that $w = v^k$.*

Proof. Since w is cyclically reduced and has nonabelian centralizer, its support is contained in a subjoin of Γ. Since Γ has no triangles and no squares, every subjoin of Γ is contained in the star of a vertex of Γ. If $v \in V(\Gamma)$, we have $\langle St(v) \rangle = F_n \times \mathbb{Z}$ for some $n \geq 2$, since Γ has no degree one vertices. The only elements of $\langle St(v) \rangle$ with nonabelian centralizer are multiples of the star vertex. □

We noted in our discussion of ping–pong a classification of right-angled Artin groups that act discretely and by isometries on hyperbolic spaces. A consequence of Mostow rigidity applied to dominated transvections is the following:

Corollary 2.20. *Let Γ be a finite simplicial graph and let $n \geq 3$. Then $A(\Gamma)$ is not the fundamental group of a finite volume hyperbolic n–manifold.*

2.7. **Two–generated subgroups.** In [2], Baudisch proves the following result:

Theorem 2.21. *Let $G = \langle g_1, g_2 \rangle < A(\Gamma)$ be a two–generated subgroup. Then G is either free or abelian.*

The proof of this useful result using combinatorial group theory is rather difficult, and we will not give it here. The ping–pong lemma for right-angled Artin groups does not seem to be of any help. We will give a more geometric proof.

Let $v \in V(\Gamma)$ be a vertex of Γ. We think of v as a one–cell of $S(\Gamma)$. The one–cell v determines a *dual hypersurface* T_v. Thinking of v as a unit vector in the x_i direction within the unit cube in $\mathbb{R}^n$, we define T_v to be the image of the hyperplane $x_i = 1/2$ inside of $S(\Gamma)$. Observe that T_v is a union of tori inside of $S(\Gamma)$, and that T_v is nonseparating inside of $S(\Gamma)$, but that T_v has a bi–collared neighborhood and is therefore locally separating. Indeed, let γ be an oriented loop in $S(\Gamma)$ based at the zero–cell of $S(\Gamma)$. Assume that γ is transverse to T_v, so that $\gamma \cap T_v$ is a finite set of points. These points all have natural orientations, coming from the fact that γ is oriented and T_v is locally separating. It follows that oriented intersection number with T_v is a homomorphism

$$\phi_v : A(\Gamma) = \pi_1(S(\Gamma)) \to \mathbb{Z}.$$

Since the loop determined by v itself takes the value 1 under ϕ_v, it follows that ϕ_v determines a nonzero integral cohomology class of $A(\Gamma)$. In particular, T_v is nonseparating in $S(\Gamma)$.

Before proving Theorem 2.21, we need two facts. The first is that a finitely generated free group is *Hopfian*, which is to say any surjection from a finitely generated free group to itself is an isomorphism. The second will be a method of constructing surjections from subgroups of $A(\Gamma)$ to free groups.

Recall that a group is called *residually finite* if the intersection of all finite index subgroups is trivial. A group is called *residually p* if the intersection of all normal subgroups of finite p–power index is trivial.

Lemma 2.22. *A finitely generated, residually finite group is Hopfian.*

Proof. Let G be finitely generated and residually finite and let $\phi : G \to G$ be a surjection with kernel K. Since G is finitely generated, there are finitely many subgroups $\{G_1, \ldots, G_k\}$ of any finite index n. Pulling back $\{G_1, \ldots, G_k\}$ under ϕ gives the same collection of subgroups $\{G_1, \ldots, G_k\}$. However, it follows that K is contained in each of these subgroups, whence K is contained in every finite index subgroup of G. Therefore K is trivial. □

The previous lemma is false if G is not finitely generated. For instance, an infinitely generated free group is a proper quotient of itself.

Lemma 2.23. *A finitely generated free group F is residually p for every prime p.*

Proof. We have that F is the fundamental group of a finite wedge of circles $X = X_0$, which we equip with the graph metric. For $i \geq 1$, we will write X_i for the regular connected covering space of X_{i-1} with deck group $H_1(X_{i-1}, \mathbb{Z}/p\mathbb{Z})$.

Let Y be a graph with the graph metric and let γ be an unbased loop without backtracking. If γ has minimal length among all nontrivial loops in Y then γ is simple in that it visits each vertex of Y at most once. Furthermore, any simple loop in Y represents a nontrivial element in $H_1(Y, \mathbb{Z})$.

It follows that the shortest loop in X_0 which lifts to X_i has length at least $i+1$. It follows that as i tends to infinity, the injectivity radius of X_i tends to infinity. The conclusion of the lemma follows. □

For the following lemma, let $X \to S(\Gamma)$ be a covering space and let T, T' be components of the preimage of dual hypersurfaces T_v and $T_{v'}$ in $S(\Gamma)$.

Lemma 2.24. *Suppose $T \cup T'$ does not separate X and that $T \cap T' = \emptyset$. Then oriented intersection with $T \cup T'$ determines a surjective homomorphism $\phi : \pi_1(X) \to F_2$.*

Proof. Let $F_2 = \langle a, b\rangle$ and let $\gamma \subset X$ be a based oriented loop which intersects both T and T' transversely. Since T and T' are disjoint, we may take small tubular neighborhoods $N(T) \cong T \times [-1, 1]$ and $N(T') \cong T' \times [-1, 1]$ which are disjoint. We define $\phi(\gamma)$ by starting at the basepoint of γ and following it around. Every time γ encounters T with a positive orientation, we record and a, and every time γ encounters T' with a positive orientation, we record a b (and similarly with a^{-1} and b^{-1} for negatively oriented intersections). This determines a homomorphism $\pi_1(X) \to F_2$. Indeed, we can collapse $N(T)$ and $N(T')$ to intervals so that topologically, ϕ is induced by a map $X \to S^1 \vee S^1$. Since both T and T' do not separate X, we have that ϕ is nontrivial. Since T does not separate $X \setminus T'$ (and similarly after switching the roles of T and T'), we have that the image in F_2 is not cyclic. □

Sketch of proof of Theorem 2.21. Assume that G is not abelian. We will find a finite index subgroup $H < A(\Gamma)$ which contains G and a surjection $\phi : H \to F_2$ such that $\phi(G)$ is not abelian, whence it will follow that $G \cong F_2$ by the Hopficity of free groups.

We may assume by induction on the number of vertices of Γ that $supp(g_1) \cup supp(g_2) = \Gamma$. By the Kurosh subgroup theorem, we may assume that Γ is a connected graph. We can also assume that Γ does not split as a nontrivial join. Indeed, if $\Gamma \cong \Gamma_1 * \Gamma_2$ for two nonempty graphs Γ_1 and Γ_2, we have $A(\Gamma) \cong A(\Gamma_1) \times A(\Gamma_2)$, and we obtain projection maps p_1 and p_2 to the two direct factors. If neither of these maps is injective when restricted to G then G splits as a direct product and is therefore abelian.

If $X \to S(\Gamma)$ is a finite cover to which a particular loop γ lifts and if $\widetilde{T}$ is a component of the preimage of a dual hypersurface T_v, then the number of intersections of any lift of γ and $\widetilde{T}$ in X is at most the number of intersections between γ and T_v. If the cover X is chosen carefully then the intersection number in X will be strictly smaller.

Observe that in $A(\Gamma)^{ab}$, the subgroup spanned by the images of g_1 and g_2 has rank at most two, so there is always a homomorphism $\phi : A(\Gamma) \to \mathbb{Z}$ whose kernel contains G. Furthermore, such a ϕ can be written as a linear combination of oriented intersection numbers with dual hypersurfaces in $S(\Gamma)$. The homomorphism ϕ can be combined with reduction modulo n in order to give a finite cover X of $S(\Gamma)$.

After taking finitely many covers of this form, one can use the fact that g_1 and g_2 do not generate a cyclic group in order to find a finite cover X of $S(\Gamma)$ with $G < \pi_1(X)$ and two disjoint dual hypersurfaces $T, T' \subset X$ whose union does not separate X and such that the oriented intersection between g_1 and T is different from the oriented intersection between g_2 and T. This furnishes a surjective map from G to F_2. □

2.8. **Residual finiteness.** Using intersection theory with dual hypersurfaces, we can easily show that right-angled Artin groups are residually finite. Thus in a precise sense, we can show that right-angled Artin groups contain "lots" of finite index subgroups.

Theorem 2.25. *For every prime p, the right-angled Artin group $A(\Gamma)$ is residually p.*

Proof. Let $S(\Gamma) = X_0$ be the Salvetti complex. We define the cover $X_i \to X_{i_1}$ to be a regular connected cover with deck group $H_1(X_{i-1}, \mathbb{Z}/p\mathbb{Z})$.

A basis for the integral cohomology of X_i is given by oriented intersection number with components of preimages of dual hypersurfaces in $S(\Gamma)$ (which we also call dual hypersurfaces). A loop $\gamma \subset X_i$ represents a nontrivial homotopy class if and only if it intersects such a dual hypersurface at least once. So, if $\gamma \subset X_i$ represents a nontrivial homotopy class, either γ does not lift to X_{i+1}, or the total intersection number of any lift of γ with a dual hypersurface in X_{i+1} decreases. Since the total number of dual hypersurfaces which γ and any of its lifts intersects is bounded by the length of γ as a word in $A(\Gamma)$, we obtain the result. □

Corollary 2.26. *Right-angled Artin groups are Hopfian.*

3. Combinatorial Morse Theory

In this section we will develop combinatorial Morse theory and use it to produce some rather exotic subgroups of right-angled Artin groups, as well as to illustrate its usefulness in combinatorial group theory. Recall that in differential topology, one studies Morse functions on manifolds. A Morse function on a manifold M is a smooth map

$$f : M \to \mathbb{R}$$

whose critical values form a discrete subset of $\mathbb{R}$ and whose critical points are non–degenerate. Outside of the critical points, M is locally a product, and various product regions meet in certain limited types of singularities at critical points. Such a function allows one to, among other things, reconstruct M as a union of local product regions and is thus useful for studying the topology of M.

If one imagines a manifold hanging from one's hand by a string, the height function provides a good intuition for a Morse function. A standard fact about Morse functions if that, at least when M is compact, the set of Morse functions is open in the function space $C^\infty(M, \mathbb{R})$. Therefore, Morse functions abound, even though it may be difficult to produce an explicit one.

Combinatorial Morse theory allows one to study the structure of affine complexes. As in the smooth case, an affine complex can be decomposed as a union of product regions which meet in a limited variety of "coning off" operations at the "critical" points of the Morse function.

As in the smooth case, if one imagines an affine complex hanging from one's hand by a string, the height function provides a good intuition for a Morse function. Let K be an affine complex. The majority of combinatorial Morse functions arise from homomorphisms of groups $\pi_1(K) \to \mathbb{Z}$, viewed topologically as an affine map $K \to S^1$, and then lifting to the universal covering space $\widetilde{K} \to \mathbb{R}$.

We will draw inspiration from N. Brady's notes in [7].

3.1. **Basic definitions.** An *affine cell complex* is a finite dimensional (though not necessarily finite) cell complex X with a piecewise affine structure. That is to say, each cell $E \subset X$ is equipped with a characteristic map

$$\chi_E : P_E \to E$$

which is a homeomorphism from a convex polyhedral cell in $\mathbb{R}^m$ to E, and such that the restriction of χ_E to a face of P_E agrees with the characteristic function of a lower–dimensional cell of X, possibly composed with an affine homeomorphism of $\mathbb{R}^m$.

A *Morse function* on an affine cell complex is a function $f : X \to \mathbb{R}$ whose restriction to each cell $E \subset X$ allows the composition $f \circ \chi_E$ to extend to an affine map $\mathbb{R}^m \to \mathbb{R}$ which is non–constant on each cell of positive dimension, and for which the 0–skeleton of X has discrete image in $\mathbb{R}$. A *circle–valued Morse function* is a map $f : X \to S^1$ whose lift to the universal covers

$$\widetilde{f} : \widetilde{X} \to \mathbb{R}$$

is a Morse function. For the sake of analogous terminology, we will call images of the 0–skeleton of X in $\mathbb{R}$ under f the *critical values* of f.

If $v \in P_E \subset \mathbb{R}^m$ is a vertex of a convex polyhedral cell, the *link* of v in P_E is written $lk_{P_E}(v)$ and is defined to be the set of inward–pointing (into P_E) unit tangent vectors of $\mathbb{R}^m$ at v. The contribution of E to the link of v in X is $\chi_E(lk_{P_E}(v))$. The link of v in X is

$$lk_X(v) = \bigcup_{v \in E} \chi_E(lk_{P_E}(v)).$$

A Morse function $f : X \to \mathbb{R}$ gives rise to two important subsets of $lk_X(v)$, namely the *ascending* and *descending* links of v. The subset $lk_X(v) \cap E$ is in the ascending link of v if f achieves its minimum on E at v. Similarly, $lk_X(v) \cap E$ is in the descending link of v if f achieves its maximum on E at v. Observe that the ascending link of v together with the descending link of v may be properly contained in the link of v, so that the ascending and descending links of v may not partition $lk_X(v)$.

3.2. **The Morse Lemma.** In differential topology, the Morse Lemma says that if $f : M \to \mathbb{R}$ is a Morse function of a manifold M allows one to decompose M as a union of product regions meeting at certain singularities. If $t \in \mathbb{R}$ is a regular value of f, then t lies in an open neighborhood $(a_t, b_t) \subset \mathbb{R}$ whose endpoints (if they are finite) are critical values of f and such that (a_t, b_t) contains no further critical values of f. The inverse function theorem implies that for every $s \in (a_t, b_t)$, there is a diffeomorphism $f^{-1}(s) \cong f^{-1}(t)$. Since the set of critical values of f are discrete, we can decompose $\mathbb{R}$ as a countable union of disjoint open intervals consisting of regular values of f, together with a countable set of critical values of f.

Let x be a critical value of f and let $(x - \epsilon, x + \epsilon)$ be a neighborhood of x which contains no other critical values of f. In general, $f^{-1}(x - \epsilon/2)$ and $f^{-1}(x + \epsilon/2)$ will not be diffeomorphic to each other. However, they fit together as parts of a singularity whose type is determined by the signature of the Hessian of f at x.

The combinatorial Morse Lemma is the analogous setup in the affine cell complex case.

Theorem 3.1. *Let $f : X \to \mathbb{R}$ be a Morse function. Suppose that $I \subset I' \subset \mathbb{R}$ are closed intervals with $\inf I = \inf I'$ such that $I' \setminus I$ contains exactly one critical value x of f. Then $f^{-1}(I')$ is homotopy equivalent to $f^{-1}(I)$ with the descending links of each vertex $v \in f^{-1}(x)$ coned off.*

Analogously, suppose that $I \subset I' \subset \mathbb{R}$ are closed intervals with $\sup I = \sup I'$ such that $I' \setminus I$ contains exactly one critical value x of f. Then $f^{-1}(I')$ is homotopy equivalent to $f^{-1}(I)$ with the ascending links of each vertex $v \in f^{-1}(x)$ coned off.

We will sketch the proof of Theorem 3.1 as it is given in [4]

Sketch of the proof of Theorem 3.1. If $Y \subset \mathbb{R}$, we will write $X_Y = f^{-1}(Y) \subset X$. The first observation is that if $I \subset I' \subset \mathbb{R}$ and I' contains no critical values of f, then X_I and $X_{I'}$ are homotopy equivalent. A straightforward argument gives a strong deformation retract from X_I to $X_{I'}$. Similarly, if $I' \setminus I$ contains a critical value then one can construct a deformation retract of X_I to $X_{I'}$ with the appropriate cones attached. □

3.3. **Free–by–cyclic groups.** The first application of combinatorial Morse theory we will give is the following result of J. Howie, the proof of which follows [7]:

Proposition 3.2. *Let X be a 2–complex and let $f : X \to S^1$ be a Morse function. Suppose that all ascending and descending links are trees. Then X is aspherical and $\pi_1(X)$ is free–by–cyclic.*

Proof. The complex X has an affine structure with respect to which f is affine. The preimage of a generic point in S^1 is a subgraph G of X. By the discreteness of the critical value set of f, we can view X as a graph of spaces whose underlying space is a circle. That is, X is a finite union of spaces which are homeomorphic to graphs cross intervals (the edge spaces) which meet along spaces given by collapsing subtrees of the graphs (the vertex spaces). Since trees are contractible, collapsing a tree is a homotopy equivalence, so that the vertex and edge spaces are homotopy equivalent to each other. It follows that X is homotopy equivalent to a graph bundle over the circle. It follows that $\pi_1(X)$ fits into a short exact sequence of the form

$$1 \to \pi_1(G) \to \pi_1(X) \to \mathbb{Z} \to 1,$$

so that $\pi_1(X)$ is free–by–cyclic. □

To illustrate Howie's result in more detail, we will consider an example. Consider the group

$$H = \langle a, b \mid aba = bab \rangle.$$

The reader may recognize this presentation as the braid group on three strands. The presentation 2–complex X has exactly one 0–cell v, two 1–cells labelled a and b, and a single 2–cell which can be taken to be a regular hexagonal cell in $\mathbb{R}^2$ with edges cyclically labelled $a, b, a, b^{-1}, a^{-1}, b^{-1}$. Instead of labeling edges a or a^{-1}, one can simply choose the opposite orientation on the edge labelled a^{-1}. We can now compute the link of v in X. The link consists of four vertices which, following Brady's notation, we label a^+, a^-, b^+, b^- according to whether the labelled edges travel away from or towards a particular vertex of the hexagonal cell, respectively. One checks that the link of v is a four–cycle whose edges are cyclically labelled a^+, b^-, a^-, b^+.

Observe that the abelianization of H is isomorphic to $\mathbb{Z}$, and the relation $aba = bab$ just becomes $a = b$. Choose the nontrivial homomorphism from H to $\mathbb{Z}$ which sends a to the generator $1 \in \mathbb{Z}$. This homomorphism gives rise to an infinite cyclic cover $\widetilde{X}$ of X.

Consider a lift of the hexagonal cell of X to $\widetilde{X}$. Drawing the cell on a page with corners towards the top and the bottom and the other four vertices in two pairs at equal heights (so that the Morse function f is just the height function on the cell), we can use Howie's result to show that $H = \pi_1(X)$ is free–by–cyclic. Note that the ascending link of v is the segment connecting a^+ to b^+, and the descending link is the segment connecting a^- to b^-. In particular, the ascending and descending links of v are trees.

The explicit nature of this example allows one to carry out a more detailed analysis. Observe that there are exactly two homeomorphism types of graphs of the form $f^{-1}(t)$, which is to say when t is a critical value of f and when it is a regular value. If t is a regular value, $f^{-1}(t)$ is a graph with two vertices and three edges connecting them. If t is a critical value, $f^{-1}(t)$ is a graph with one vertex v and two edges. It follows that the rank of the kernel of the map $H \to \mathbb{Z}$ has rank two.

By the Morse Lemma, varying t through a critical value of f results in a coning off of an ascending or a descending link of v, which is a homotopy equivalence in our case. One can easily see that the homotopy equivalence between the two homeomorphism types of graphs is given by collapsing one of the edges in the regular preimage graph to a point.

Observe that the deck group $\mathbb{Z}$ acts on $\pi_1(X) = F_2$. In our case, one can compute the induced automorphism of F_2 by appropriately composing several of the homotopy equivalences of the graphs $\{f^{-1}(t)\}$ above. It is an instructive exercise to compute this automorphism.

3.4. **Exotic subgroups of right-angled Artin groups.** We will soon want to use Morse theory to systematically build families of subgroups of right-angled Artin groups which themselves are not right-angled Artin groups. Subgroups of right-angled Artin groups which are not themselves right-angled Artin groups abound in general, as is illustrated by the following result of Droms in [19]:

Theorem 3.3. *Let Γ be a finite simplicial graph and let $A(\Gamma)$ be the associated right-angled Artin group. Every finitely generated subgroup of $A(\Gamma)$ is itself a right-angled Artin group if and only if Γ has no induced squares and no induced path on four vertices.*

Proof. Suppose that Γ has an induced square. Then $A(\Gamma)$ contains a copy of $F_2 \times F_2$. It is a standard fact that $F_2 \times F_2$ contains subgroups which are finitely generated but not finitely presented, which therefore cannot be right-angled Artin groups. Alternatively, we will construct an explicit example immediately following the proof.

Suppose that Γ has an induced path on four vertices, P_4. Label the vertices linearly by x, y, z, w, and let

$$\phi : A(P_4) \to \mathbb{Z}/2\mathbb{Z}$$

be the homomorphism which sends each of the vertices to the nontrivial element of $\mathbb{Z}/2\mathbb{Z}$. Writing a, b, c, d for the elements $x^{-1}y, y^2, y^{-1}z, z^{-1}w$ respectively and K for the kernel of ϕ, the Reidemeister–Schreier rewriting process shows that K has the presentation

$$K = \langle a, b, c, d \mid ab = ba,\ bc = cb,\ bc^2d = dbc^2\rangle.$$

We claim that K is not isomorphic to any right-angled Artin group. To see this, note that $K^{ab} \cong \mathbb{Z}^4$, so that any candidate graph Γ would have to have exactly four vertices. Furthermore, Γ cannot have any triangles by a cohomological dimension argument. Observe that since K has finite index in $A(P_4)$, the Euler characteristics of these two groups must both be zero. One easily checks that the Euler characteristic of a right-angled Artin group on a connected triangle–free graph Γ differs from the Euler characteristic of Γ by exactly one. If K is a right-angled Artin group then the underlying graph must be connected, since K has one end. Indeed, K has finite index in $A(P_4)$, which has one end. It follows that Γ must be a tree. There are only two possibilities for Γ, namely P_4 and the tree $D_3 * v$ with one vertex of degree three and three vertices of degree one.

We can rule out both of these possibilities. The second possibility of $D_3 * v$ is ruled out by the fact that K has no center. This can be seen by describing K as an amalgamated free product. We have that

$$K = \langle a, b, c \mid ab = ba, bc = cb\rangle *_{bc^2=e} \langle e, d \mid ed = de\rangle.$$

Since bc^2 and a generate a free group, K cannot have a nontrivial center. To see that K is not isomorphic to $A(P_4)$, we follow Droms and pass to the universal 2–step nilpotent quotient of K and $A(P_4)$. Write

$$K_2 = K/[K, [K, K]]$$

and G_2 for the corresponding universal 2–step nilpotent quotient of $A(P_4)$. Write K_2' and G_2' for the derived subgroups of K_2 and G_2. Observe that the commutator identity

$$1 = [bc^2, d] = [b, d][c, d]^2$$

in K_2 implies that the image of b is central in $K_2/(K_2')^2$. Furthermore, one sees that the image of b is itself not a square in $K^{ab} \cong \mathbb{Z}^4$. A straightforward computation shows that if $g \in G_2$ is central modulo $(G_2')^2$ then g is itself a square in $A(P_4)^{ab}$, a contradiction.

We now establish the converse. Note that if Γ is disconnected then $A(\Gamma)$ decomposes as a free product of the right-angled Artin groups on the components. By the Kurosh Subgroup Theorem, a finitely generated subgroup of a free product is a free product of conjugates of subgroups of the free product factors, possibly together with a free group free factor. We may therefore suppose that Γ is connected.

It is an elementary exercise in graph theory to show that if Γ is a finite connected graph with no induced copy of P_4 or of a square then Γ decomposes as a join $\Gamma = \Gamma_0 * v$ for some vertex $v \in V(\Gamma)$.

We can establish the conclusion now by induction on the number of vertices of Γ. The conclusion is clear when Γ has exactly one vertex. Write $\Gamma = \Gamma_0 * v$, and let $H < A(\Gamma)$ be a finitely generated subgroup. Write $p : A(\Gamma) \to A(\Gamma_0)$ for the projection map given by killing the vertex v. The group $p(H) < A(\Gamma_0)$ is a right-angled Artin group by the inductive hypothesis. So, we may write $p(H) = A(\Lambda)$. For each $\lambda \in V(\Lambda)$, choose a preimage $i(\lambda) \in H$. Any two preimages of λ differ by a multiple of $\langle v \rangle$. Since v is central in $A(\Gamma)$, we have that i defines a homomorphism which splits p. It follows that H is either isomorphic to $A(\Lambda)$ or $A(\Lambda) \times \mathbb{Z}$, which is what we set out to prove. □

Before returning to Morse theory, we will note one an exotic subgroup of $F_2 \times F_2$, as promised. This example will be important in the sequel. Write $F_2 \times F_2 = \langle x, y \rangle \times \langle w, z \rangle$, and let $g = yz$. The following proposition appears in [32]:

Proposition 3.4. *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 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 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. □

3.5. **Homological algebra.** We now develop some of the tools we will need in order to analyze Bestvina–Brady kernels and to formulate and prove the relevant Morse–theoretic criterion. We follow the exposition in [4]. We will write H for a discrete group and R for a commutative ring with nonzero identity. As is usual in homological algebra, we consider R as a trivial module over the group ring $R[H]$.

The group H is of *type* $FP_n(R)$ if there exists a resolution of R by finitely generated projective $R[H]$–modules of length n. That is to say, there is an exact sequence of the form

$$P_n \to P_{n-1} \to \cdots \to P_0 \to R,$$

where for each i the module P_i is finitely generated and projective. A group is $FP(R)$ is there exists such a finite resolution of the form

$$0 \to P_n \to P_{n-1} \to \cdots \to P_0 \to R.$$

The group H is of *type* F_n if there exists a $K(H, 1)$ with a finite n–skeleton.

It would be useful to have an algebraic criterion for determining when a group is $FP_n(R)$ but not $FP_{n+1}(R)$. The following result gives us just that:

Proposition 3.5. *Suppose there exists a projective resolution of the $R[H]$ module R of the form*

$$0 \to M_n \to P_n \to \cdots \to P_0 \to R \to 0,$$

where each P_i is finitely generated and projective, and where M_n is not finitely generated. Then H is of type FP_n but not of type FP_{n+1}.

This proposition is a consequence of a general version of the well–known Schanuel's Lemma, which can be found in [10] for instance. For the convenience of the reader, we recall the statement with proof in the short exact sequence case:

Lemma 3.6 (Schanuel's Lemma). *Let*

$$0 \to K_1 \to P_1 \to M \to 0$$

and

$$0 \to K_2 \to P_2 \to M \to 0$$

be short exact sequences of modules over R, with P_1 and P_2 projective. Then $P_2 \oplus K_1 \cong P_1 \oplus K_2$.

Observe that this version of Schanuel's lemma can be immediately applied in the case of Proposition 3.5 when $n = 0$. Indeed, suppose there exists a resolution of R of the form

$$0 \to M_0 \to P_0 \to R \to 0$$

where M_0 is not finitely generated, but that there exists a projective resolution

$$P_1' \to P_0' \to R \to 0.$$

Then if K denotes the kernel of the map $P_1' \to P_0'$, we have a resolution

$$0 \to P_1'/K \to P_0' \to R \to 0$$

by finitely generated modules. Schanuel's Lemma implies that $M_0 \oplus P_0'$ is finitely generated, a contradiction.

Proof of Schanuel's Lemma. Let ϕ_1 and ϕ_2 denote the two maps $P_1 \to M$ and $P_2 \to M$. Let $S \subset P_1 \oplus P_2$ be the associated fiber product. There are natural maps $S \to M$ for each i given by projection onto one of the factors and composing with ϕ_i, which are both surjective since each ϕ_i is surjective. This implies that the projection map of S onto either factor is also surjective. Write π_1 and π_2 for the two projection maps. An easy computation shows that $\ker \pi_1 = K_2$ and $\ker \pi_2 = K_1$. In particular, we have a pair of exact sequences of the form

$$0 \to K_i \to S \to P_i \to 0$$

which both split, since P_i is projective. The claim follows. □

A group H is of *type* $FH_n(R)$ if it acts freely, faithfully, properly discontinuously, cellularly and cocompactly on a cell complex X whose reduced homology satisfies $\widetilde{H}_i(X, R) = 0$ for all $i \leq n-1$.

A group H is of *type* $FH(R)$ if it acts freely, faithfully, properly discontinuously, cellularly and cocompactly on an R–acyclic cell complex X.

The proof of the following lemma is straightforward and left as an exercise to the reader.

Lemma 3.7. *Let H be a group and a R a commutative ring as above.*

1. *If $H \in FH_n(R)$ then $H \in FP_n(R)$.*
2. *If $H \in FH(R)$ then $H \in FP(R)$.*
3. *If H acts freely, faithfully, properly discontinuously, cellularly and cocompactly on an $(n-1)$–connected cell complex X then H is of type F_n.*

There is a further finiteness property called *type* $FL_n(R)$ which is defined the same way as type $FP_n(R)$ with the word "projective" replaced by the word "free". The finiteness properties are related as follows:

$$F_n \to FH_n \to FL_n \to FP_n.$$

We obtain a topological criterion for showing a group is $FP_n(R)$ but not $FP_{n+1}(R)$ from the previous lemma:

Corollary 3.8. *Suppose that H acts freely, faithfully, properly discontinuously, cellularly and cocompactly on a cell complex X for which $\widetilde{H}_i(X, R) = 0$ for $0 \leq i \leq n-1$ and for which $\widetilde{H}_n(X, R)$ is not a finitely generated $R[H]$–module. Then H is of type $FP_n(R)$ but not of type $FP_{n+1}(R)$.*

3.6. **Kernel subgroups and their finiteness properties.** We are now ready to use Morse theory to study finiteness properties of groups. We first briefly revisit the Morse Lemma:

Proposition 3.9 ([4])**.** *Let X be an affine cell complex and let $f : X \to \mathbb{R}$ be a Morse function. Suppose that $J \subset J' \subset \mathbb{R}$ are intervals, and let X_J and $X_{J'}$ denote $f^{-1}(J)$ and $f^{-1}(J')$ respectively.*

(1) *If each descending and ascending link is homologically n–connected then the inclusion $X_J \subset X_{J'}$ induces an isomorphism on the reduced homology groups $\widetilde{H}_i$ for $i \leq n$ and a surjection on $\widetilde{H}_{n+1}$.*
(2) *If each descending and ascending link is simply connected then the inclusion $X_J \subset X_{J'}$ is an isomorphism on fundamental groups.*
(3) *If each descending and ascending link is connected then the inclusion $X_J \subset X_{J'}$ is a surjection on fundamental groups.*

Proof. This follows from a straightforward application of the Mayer–Vietoris sequence and the Seifert–Van Kampen Theorem. □

Let G act freely, faithfully, properly discontinuously, cellularly and cocompactly on a contractible affine cell complex X by affine homeomorphisms. Let $\phi : G \to \mathbb{Z}$ be a surjective map with kernel H and let $f : X \to \mathbb{R}$ be a Morse function which is ϕ–equivariant.

Theorem 3.10 ([4], cf. [7])**.** *With the setup above, we have the following:*

(1) *Suppose each ascending and descending link is homologically n–connected. Then $H \in FH_{n+1}(R)$.*
(2) *Suppose each ascending and descending link is R–acyclic. Then $H \in FH(R)$.*
(3) *Suppose each ascending and descending link is simply connected. Then H is finitely presented.*

Proof. By Proposition 3.9, we have that for $t < s$, the inclusion $X_{(-\infty,t]} \to X_{(-\infty,s]}$ induces isomorphisms on $\widetilde{H}_i$ for all $i \leq n$ and a surjection on $\widetilde{H}_{n+1}$. Writing

$$X = \bigcup_{r \in \mathbb{Z}} X_{(-\infty,r]}$$

and using the fact that X is acyclic, we have that $\widetilde{H}_i(X_{(-\infty,t]})$ vanishes for all $i \leq n$ and all t. Writing

$$X = X_{(-\infty,t]} \cup X_{[t,\infty)},$$

we see that X_t is homologically n–connected. The action of H on X_t is free and cocompact, so that H is of type $FH_{n+1}(R)$. This establishes item (1). Item (2) follows similarly. For item (3), we have that X_t is connected for each t, and Proposition 3.9 implies that the inclusion $X_t \to X$ induces an isomorphism of fundamental groups. It follows that H is finitely presented. □

The following criterion can be used to give another criterion for showing that a group is FH_n but not FH_{n+1} and results more or less from Schanuel's lemma:

Theorem 3.11. *Let X and H be as above. Suppose X is $(n+1)$–dimensional and that the ascending and descending links of each vertex are homologically $(n-1)$–connected. Suppose furthermore that $\widetilde{H}_n$ does not vanish on the descending or ascending link of some vertex of X. Then H is of type FH_n but not FH_{n+1}.*

The proof is left to the reader.

3.7. **The free group.** It is well–known that if F_k is a nonabelian free group and $H < F_k$ is any nontrivial normal subgroup then H is finitely generated if and only if H has finite index in F_k. The standard proof goes something like this: if H has finite index then it is clearly finitely generated. Conversely, view F_k as a wedge R of k circles. The subgroup H gives rise to a regular covering space $\widetilde{R}$ of R, together with a deck group action by the infinite group F_k/H. Since H is nontrivial, there is a simple closed loop in $\widetilde{R}$. By the proper discontinuity of the action of F_k/H, there are infinitely many disjoint closed loops in $\widetilde{R}$, which are homologically independent of each other. Thus, H cannot be finitely generated since its abelianization has infinite rank.

The first case of the Bestvina–Brady theory is the Morse–theoretic analysis of the kernel of the map $F_2 = \langle a, b\rangle \to \mathbb{Z}$ given by sending both a and b to the generator $1 \in \mathbb{Z}$.

Proposition 3.12. *Let $H < F_2 = \langle a, b\rangle$ be the kernel of the map $\phi : F_2 \to \mathbb{Z}$ given by sending $a, b \to 1 \in \mathbb{Z}$. Then H is not finitely generated.*

Proof. The map ϕ can be topologically realized as an affine map $S^1 \vee S^1 \to S^1$ with their usual cell structures, given by sending the wedge point to the 1–cell of S^1 and by wrapping each copy of S^1 in the source around the circle once. Lifting this map to the universal cover X of $S^1 \vee S^1$, we obtain a Morse function from the regular 4–valent tree to $\mathbb{R}$. The ascending and descending link of each vertex consists of a copy of the 0–sphere.

Let t be a critical point of f. Then $f^{-1}(t) = X_t$ is an infinite collection of 0–cells. The group H acts freely on this set, and the quotient is a single vertex. Thus, there is a bijective correspondence between elements of H and points in X_t. It follows that if $H_0(X_t, \mathbb{Z})$ is not finitely generated as a $\mathbb{Z}[H]$–module then H is not finitely generated, since if H were finitely generated we could attach finitely many 1–cells to X_t in an H–equivariant fashion in order to get a copy of the Cayley graph of H.

Suppose that we could attach 1–cells $\{e_1, \ldots, e_m\}$ to X_t in an H–equivariant way to get a connected graph. Since X is connected, we may attach these cells within X. The H–equivariance of the attaching maps implies that $X_t \cup H \cdot \{e_1, \ldots, e_m\} \subset X_I$ for some finite interval $I \subset \mathbb{R}$ which contains t.

The Morse Lemma says that the inclusion $X_t \to X_I$ is given by taking $X_t \times I$ and coning off the appropriate 0–spheres at the preimages of critical values of f. Thus if $X_t \cup H \cdot \{e_1, \ldots, e_m\} \subset X_I$ is connected then so is X_I. Let v be a vertex above X_I. Then there are two edges descending from v which we may extend to paths into X_I. Since X_I is connected, we may complete these two paths in X_I to a closed loop in X. Since X is a tree, this is impossible. □

3.8. **Products of free groups.** We will now generalize the Morse theoretic ideas for free groups to products of free groups. The kernel subgroups we get this time will be finitely generated but not finitely presented. We first give a direct proof of the finite generation of certain kernel subgroups.

Proposition 3.13. *Let K_1 and K_2 be the kernels of the maps*

$$\phi_1 : F_2 \times \mathbb{Z} \to \mathbb{Z}$$

and

$$\phi_2 : F_2 \times F_2 \to \mathbb{Z},$$

viewed as right-angled Artin groups, given by sending each vertex generator to the element $1 \in \mathbb{Z}$. Then $K_1 \cong F_2$ and K_2 is finitely generated.

The abstract isomorphism type of K_2 is difficult to describe since, as we will see soon, K_2 is not finitely presented.

Proof of Proposition 3.13. Write $F_2 \times \mathbb{Z} \cong \langle a, b\rangle \times \langle t\rangle$. We claim that K_1 is generated by at^{-1} and bt^{-1}. Certainly both of these elements are in the kernel of the map ϕ_1. In general, an element $w(a, b)t^n = F_2 \times \mathbb{Z}$, where $w = w(a, b)$ is a word in a and b, is contained in the kernel of ϕ_1 if and only if the sum of the a–exponents and b–exponents in w is $-n$. But any such word can be written in the generators at^{-1} and bt^{-1}, treating the first as a and the second as b. Since t is central, we can move all the occurrences of t to the right in order to get a word of the form $w(a,b)t^n$. Under the projection map $F_2 \times \mathbb{Z} \to F_2$ given by sending t to the identity, the elements at^{-1} and bt^{-1} get sent to a and b respectively. Since finitely generated free groups are Hopfian, it follows that $\langle at^{-1}, bt^{-1}\rangle \cong F_2$ (alternatively one can appeal to the classification of two–generators subgroups of right-angled Artin groups, but that is a rather powerful result to apply here).

Let $F_2 \times F_2$ be presented as $\langle a,b\rangle \times \langle c,d\rangle$. We now claim that K_2 is generated by

$$\{ac^{-1}, bc^{-1}, bd^{-1}\}.$$

Clearly each of these elements lies in the kernel of ϕ_2. Any elements of K_2 can be written in the form $w(a,b)w'(c,d)$, where w and w' are words in the free groups and $\{a,b\}$ and $\{c,d\}$ respectively, and where the exponent sums of a, b, c and d add up to zero.

Write $w(a,b)$ as a word in ac^{-1} and bc^{-1}, moving the occurrences of c to the far right, to get a word of the form $w(a,b)c^n$. Now, write the word $c^{-n}w'(c,d)$ as a word in cb^{-1} and db^{-1}, moving the occurrences of b to the far right. Observe that all the occurrences of b cancel. Indeed, the exponent sums of a and b in $w(a,b)$ add up to $-n$, so that the exponent sums of c and d in $w'(c,d)$ add up to n. Thus, the total exponent sum in $c^{-n}w'(c,d)$ is zero, whence all the occurrences of b cancel. Concatenating these two words gives us $w(a,b)w'(c,d)$. □

Whereas K_1 is just a free group, K_2 defies such a simple description.

Theorem 3.14. *The group K_2 above is not finitely presented.*

Proof. We will give a Morse–theoretic proof. $F_2 \times F_2$ is the fundamental group of the space

$$X = (S^1 \vee S^1) \times (S^1 \vee S^1).$$

An affine cell structure can be given by the product of the usual affine structures on the two copies of $S^1 \vee S^1$. The one–skeleton consists of four one–cells meeting at the single zero–cell, and the two–skeleton is built out of four two–cells.

The map $f : X \to S^1$ maps the zero–cell to the zero–cell of S^1 with the usual cell structure, and is a homeomorphism on each one–cell. The map f is defined in the obvious way on the two–cells of X, so that the preimage of a point in S^1 outside of the zero–skeleton is a pair of "diagonals" in each two–cell. Note that f induces the map ϕ_2 of Proposition 3.13. Note that the link of the zero–cell of X consists of four points, namely the graph D_4, so that the link of the zero–cell of X is a join of two sets consisting of four disconnected points, i.e. $D_4 * D_4$.

The universal cover $\widetilde{X}$ is a product of two regular 4–valent trees. The Morse function f lifts to a real–valued Morse function $\widetilde{f}$, and the preimage of an integer point t is the union of the diagonals connecting vertices of the two–cells which lie at level t. Note that in each tree in the product, the ascending and descending links are just pairs of vertices, so that the ascending and descending links of the zero–cell of $\widetilde{X}$ are both copies of the graph $D_2 * D_2$, which is homeomorphic to S^1. We have that K_2 acts on $\widetilde{f}^{-1}(t)$ with compact quotient, and that $\widetilde{f}^{-1}(t)$ is connected.

First, if$\widetilde{f}^{-1}(t)$ is connected then H is finitely generated. Indeed, if $\widetilde{f}^{-1}(t)$ is connected, it follows that K_2 is a quotient of the fundamental group of the preimage of the zero–cell of S^1 under f, so that K_2 is finitely generated (the reader may check that this argument shows K_2 is a quotient of a free group of rank four (provided $\widetilde{f}^{-1}(t)$ is connected), so that the generating set produced in the proof of Proposition 3.13 is sharper).

Secondly, we claim that if $H_1(\widetilde{f}^{-1}(t), \mathbb{Z})$ is not finitely generated as a module over $\mathbb{Z}[K_2]$, then K_2 is not finitely presented. We write $F_2 \times F_2$ as above by

$$\langle a,b\rangle \times \langle c,d\rangle.$$

Since the quotient of $\widetilde{f}^{-1}(t)$ by K_2 has one zero–cell and four 1–cells labelled by $ac^{-1}, bc^{-1}, ad^{-1}, bd^{-1}$, we have that $\widetilde{f}^{-1}(t)$ is the Cayley graph of K_2 with respect to this generating set. If K_2 were finitely presented then we could attach finitely many 2–cells to $\widetilde{f}^{-1}(t)$ in a K_2–equivariant fashion in order to get a simply–connected Cayley complex for K_2, which implies that $H_1(\widetilde{f}^{-1}(t), \mathbb{Z})$ must be finitely generated as a module over $\mathbb{Z}[K_2]$.

We claim that $\widetilde{f}^{-1}(t)$ is indeed connected and that no such attachment of 2–cells exists. The Morse Lemma implies that as we take larger and larger preimages of intervals in $\mathbb{R}$ under f^{-1}, the only topological changes occur after passing through integer points in $\mathbb{R}$, which result in the coning off of of ascending or descending links. It follows that $\widetilde{f}^{-1}(t)$ is connected (this follows from the more general result of Bestvina–Brady described above, but can also be verified directly in a rather straightforward manner).

The proof that $H_1(\widetilde{f}^{-1}(t), \mathbb{Z})$ is not finitely generated as a module over $\mathbb{Z}[K_2]$ proceeds similarly to the case of the free group. We may attach the 2–cells in a K_2–equivariant way within $\widetilde{X}$ itself, and we may assume that all the attached 2–cells lie in $\widetilde{f}^{-1}(I)$ for some finite interval $I \subset \mathbb{R}$. We have that $\widetilde{f}^{-1}(I)$ has trivial first homology, since coning off circles cannot give rise to new homology classes (alternatively one can appeal to the more general result of Bestvina–Brady). However, as in the case of the free group, we can easily build a nontrivial 2–cycle in $\widetilde{X}$, which is impossible since $\widetilde{X}$ is contractible. □

Recall that when we wanted to show that not every finitely generated subgroup of $F_2 \times F_2$ is a right-angled Artin group in the proof of Droms' Theorem, we considered the subgroup G of

$$F_2 \times F_2 \cong \langle a, b\rangle \times \langle c, d\rangle$$

generated by $g = ac$, b and d and we showed directly that this group is not isomorphic to any right-angled Artin group. We are now in a position to show that G is not finitely presented and thus not isomorphic to a right-angled Artin group.

Observe that replacing c by c^{-1} and fixing the other generators of $F_2 \times F_2$, we get a group isomorphic to G generated by ac^{-1}, b and d.

Proposition 3.15. *The group G is isomorphic to the kernel K of the map $\phi : F_2 \times F_2 \to \mathbb{Z}$ given by $a, c \mapsto 1$ and $b, d \mapsto 0$.*

Proof. The map ϕ measures the sum of the a and c exponents of an element of $F_2 \times F_2$. A typical element of $F_2 \times F_2$ is written as $w(a, b)w'(c, d)$. Write $a(w)$ for the a–exponent of w and $c(w')$ for the c–exponent of w'. We have that $w(a, b)w'(c, d)$ is in K if and only if $a(w) - c(w') = 0$. We claim that ac^{-1}, b and d generate K.

Write $w(a, b)$, replacing each occurrence of a with ac^{-1} and moving all the occurrences of c to the right. This way, we get the word $w(a, b)c^{a(w)}$. Now write the word $c^{-a(w)}w'(c, d)$, replacing each occurrence of c with ca^{-1} and moving all the occurrences of a to the right. This way, we get the word

$$c^{-a(w)}w'(c, d)a^{a(w)+c(w')} = c^{-a(w)}w'(c, d).$$

Concatenating these two words, we get $w(a, b)w'(c, d)$, whence the claim. □

The homomorphism ϕ is not very well–adapted for a Morse–theoretic analysis of K. Indeed, the canonical choice of circle–valued functions realizing the homomorphism ϕ on

$$X = (S^1 \vee S^1) \times (S^1 \vee S^1)$$

would be constant on at least two of the one–cells of X. However, we have the following:

Proposition 3.16. *There is an isomorphism of groups $K \cong K_2$. In particular, K is not finitely presented.*

Proof. The assignments $b \mapsto ab$ and $d \mapsto cd$ extend to an automorphism of $F_2 \times F_2$, so that we may write

$$F_2 \times F_2 \cong \langle a, ab\rangle \times \langle c, cd\rangle.$$

Observe that the map ϕ maps each of these new generators of $F_2 \times F_2$ to $1 \in \mathbb{Z}$, so that we have $\ker \phi \cong K_2$. □

The reader may suspect now that the Morse–theoretic analysis of K_2 generalizes to arbitrary finite products of free groups, which it does after replacing the relevant objects by their analogues in higher dimensions. We will not work out any examples here in complete detail, and we will just note some of the relevant facts.

Consider a finite direct product

$$G = (F_2)^n = F_2 \times \cdots \times F_2,$$

and let

$$X = \bigvee_{i=1}^{n} S^1_i,$$

where we give X the product affine structure. The case where the ranks of the free groups are different is much the same. Then $G = \pi_1(X)$. The universal cover $\widetilde{X}$ is just a product of n regular 4–valent trees. We define a circle–valued Morse function $f : X \to S^1$ as before with induced kernel K_n on fundamental groups, and we build a Morse function $\widetilde{f} : \widetilde{X} \to \mathbb{R}$ by pulling back.

Again, the link of the zero–cell of $S^1 \vee S^1$ is D_4, so that the link of the zero–cell of X is just an n–fold join of D_4. As in the case of $F_2 \times F_2$, the ascending and descending links are n–fold joins of D_2, and it is a useful exercise for the reader to check that such an n–fold join is homeomorphic to the $(n-1)$–sphere. In particular, the topological change induced by varying $\widetilde{f}^{-1}(t)$ over a critical value of $\widetilde{f}$ is the coning off of a collection of $(n-1)$–spheres, which does not affect the homological $(n-1)$–connectivity and the simple–connectivity of the level sets. Since a product of trees is contractible, we have that $\widetilde{f}^{-1}(t)$ is simply connected for all t. As in the free group and $F_2 \times F_2$ example, we can glue in finitely many n–cells (within $\widetilde{X}$) in a K_n–equivariant way in order to obtain a finite interval $I \subset \mathbb{R}$ containing a critical value of $\widetilde{f}$ such that $H_{n-1}(\widetilde{f}^{-1}(I), \mathbb{Z})$ is trivial. As before, we can build a nontrivial n–cycle within X, which is impossible since X is contractible. It follows that $H_{n-1}(\widetilde{f}^{-1}(t), \mathbb{Z})$ is not finitely generated as a $\mathbb{Z}[K_n]$–module.

The fact that $\widetilde{f}^{-1}(t)$ is simply connected implies that K_n is finitely presented. Indeed, K_n acts on this complex by covering transformations with compact quotient Y, so that $\widetilde{f}^{-1}(t)$ is the universal cover of Y. It follows that $K_n = \pi_1(Y)$, so that K_n is finitely presented.

Finally, a straightforward application of Schanuel's Lemma (see Proposition 3.5) shows that if K_n were of type F_n then $H_{n-1}(\widetilde{f}^{-1}(t), \mathbb{Z})$ would have to be finitely generated as a module over $\mathbb{Z}[K_n]$.

3.9. **Bestvina–Brady subgroups of right-angled Artin groups.** We will follow [7] in our discussion here. Let Γ be a finite simplicial graph, and let K_Γ be the kernel of the map $A(\Gamma) \to \mathbb{Z}$ defined by $v \mapsto 1$ for each $v \in V(\Gamma)$. We will call K_Γ the *Bestvina–Brady subgroup* (or the *Bestvina–Brady kernel*) of $A(\Gamma)$. In order to understand the structure of K_Γ better, we need a slightly different perspective on the Salvetti complex $S(\Gamma)$.

We define the *flag complex* $F(\Gamma)$ of Γ by requiring each complete subgraph on $n+1$ vertices to span an n–simplex. If Γ has n vertices then the flag complex $F(\Gamma)$ can be realized within the unit sphere of $\mathbb{R}^n$ as follows: identify each vertex of Γ with the endpoint of a standard unit basis vector in $\mathbb{R}^n$. If $\Gamma_0 \subset \Gamma$ is a complete subgraph, we take the convex hull of the endpoints corresponding to the vertices in Γ_0. The union of these convex hulls is homeomorphic to $F(\Gamma)$. It is a straightforward exercise to show that a finite simplicial complex is homeomorphic to the flag complex of some finite simplicial graph.

Now we can build a geometric model of $S(\Gamma)$ by taking the union of cubes $C(\Gamma)$ in the unit cube in $\mathbb{R}^n$ so that the link of the origin in $C(\Gamma)$ is $F(\Gamma)$. That is, we take a 1–dimensional unit cube for each standard unit basis vector in $\mathbb{R}^n$ and declare that a collection of k of these spans a k–cube if the corresponding subgraph of Γ is complete. The Salvetti complex $S(\Gamma)$ is homeomorphic to the quotient of $C(\Gamma)$ by the usual $\mathbb{Z}^n$–action on $\mathbb{R}^n$. We leave the verification of this fact to the reader.

Let X be an arbitrary simplicial complex. We define the *sphere complex* $Sph(X)$ as follows: each k–simplex of X is replaced by a k–sphere, and if two simplices meet along a face then they meet along an appropriate sub–sphere. Precisely, let $\{v_0, \ldots, v_k\}$ span a k–simplex, which we view as the join

$$\{v_0\} * \cdots * \{v_k\}.$$

We replace each vertex v_i by a pair of disjoint vertices $\{v_i^{\pm}\}$ and replace the simplex by the k–sphere

$$\{v_0^{\pm}\} * \cdots * \{v_k^{\pm}\}.$$

Proposition 3.17. *The link of the* 0*–cell of* $S(\Gamma)$ *in the construction above is homeomorphic to* $Sph(F(\Gamma))$*.*

Proof. Consider the link of the origin in $[0,1]^n/\mathbb{Z}^n$. It is easy to check that the link is homeomorphic to an n–fold star of $\{-1, 1\}$, and that this homeomorphism is compatible with coordinate inclusions of unit cubes. □

Observe that one can easily define a circle–valued Morse function f on $S(\Gamma)$ whose induced map on fundamental groups is the Bestvina–Brady map $A(\Gamma) \to \mathbb{Z}$. The one–cubes (viewed as unit length segments along coordinate vectors in $\mathbb{R}^n$) are mapped to the one–cell of S^1, and the extension to the interior of the higher–dimensional cubes of $S(\Gamma)$ is by linear interpolation. We thus get a Morse function

$$\widetilde{f} : \widetilde{S(\Gamma)} \to \mathbb{R}.$$

Observe that the ascending and descending link of each vertex of $\widetilde{S(\Gamma)}$ is homeomorphic to $F(\Gamma)$.

The following result is now immediate in view of the homological algebra we gave discussed:

Theorem 3.18. *Let* Γ *be a finite simplicial graph. If* $F(\Gamma)$ *is homologically* n*–connected over* R *then* K_Γ *is of type* $FH_{n+1}(R)$*. If* $F(\Gamma)$ *is acyclic over* R *then* K_Γ *is of type* $FH(R)$*. If* $F(\Gamma)$ *is simply connected then* K_Γ *is finitely presented.*

We would now like to prove some of the reverse implications of the previous result, following [4]. A *sheet* in $\widetilde{S(\Gamma)}$ is a flat which is a subcomplex of the preimage of one of the standard tori of $S(\Gamma)$. If v is a vertex of $\widetilde{S(\Gamma)}$ and $A \subset \widetilde{S(\Gamma)}$ is a union of sheets, we can define the *star* of v in A as the cone on the link of v in A. Observe that this star is homeomorphic to a small neighborhood of v in A. Similarly, one can define the ascending and descending stars of v as cones of the ascending and descending stars of v in A.

Observe that there is a natural retraction r_v from the link of v in $\widetilde{S(\Gamma)}$ to the ascending (respectively descending) link of v in $\widetilde{S(\Gamma)}$ which restricts to the identity on the ascending (respectively descending) link of v. It is a useful exercise for the reader to determine the exact form of r_v. Note the r_v can be coned off to give a retraction from the star of v to the ascending (respectively descending) star of v.

The following proposition, found in [4], summarizes the properties of sheets. The proof is easy and is left as an exercise for the reader.

Proposition 3.19. *Let* $\widetilde{S(\Gamma)}$ *be as above.*

(1) *The complex* $\widetilde{S(\Gamma)}$ *is covered by sheets.*

(2) *The intersection of any collection of sheets is either empty, a vertex, or a sheet.*
(3) *All (sub)–level set of $\widetilde{f}$ restricted to a sheet are contractible, and all ascending and descending links of the restriction are single simplices.*
(4) *The retraction r_v preserves sheets containing v.*

For a vertex $v \in \widetilde{S(\Gamma)}$ and A a union of sheets, we will write $lk_0(v, A)$ and $st_0(v, A)$ (respectively $lk^0(v, A)$ and $st^0(v, A)$) for the descending (respectively ascending) link and star of v in A, respectively. For an interval $I \subset \mathbb{R}$, we will write $\widetilde{S(\Gamma)}_I$ and A_I for $\widetilde{f}^{-1}(I)$ and $\widetilde{f}^{-1}(I) \cap A$ respectively. The following lemma is found in [4]:

Lemma 3.20. *For any ring of coefficients, we have an isomorphism*

$$H_*(A, A_J) \cong \bigoplus_{v \notin A_j} H_*(st_0(v, A), lk_0(v, A)).$$

Proof. If $v \in A \setminus A_J$, a map

$$\Psi_A : H_*(A, A_J) \cong H_*(st_0(v, A), lk_0(v, A))$$

can be defined by a composition of maps

$$H_*(A, A_J) \to H_*(A, A \setminus Int(st(v, A))) \cong H_*(st(v, A), lk(v, A)) \to H_*(st_0(v, A), lk_0(v, A)),$$

where the last map is induced by the retraction r_v. Observe that Ψ_A vanishes for all but finitely many v, so that the map into the direct sum is well–defined, and that Ψ_A commutes with inclusions between unions of sheets (i.e. Ψ_A is natural).

We now claim that the map Ψ_A is an isomorphism. First, it suffices to prove the claim when A is a finite union of sheets, since homology commutes with direct limits. Then, one proceeds on the number of sheets in A. The base case is when A is empty, a single sheet, or a vertex, and is straightforward. Otherwise, suppose A is a union of $k+1$ sheets and vertices, so that $A = A' \cup S$, where S is a sheet or a vertex and A is a union of k sheets and vertices. We have that $A' \cap S$ is a union of at most k sheets and vertices, by the previous proposition. The fact that Ψ_A is an isomorphism follows from the Mayer–Vietoris sequence, the naturality of Ψ_A, and the 5–lemma. □

As a corollary of the lemma, one can easily show the following now:

Corollary 3.21. *For any ring R, here are $R[K_\Gamma]$–isomorphisms*

$$\widetilde{H}_*(\widetilde{S(\Gamma)}_{(-\infty,t]}) \cong \bigoplus_{v \in (t,\infty)} \widetilde{H}_*(F(\Gamma))$$

and

$$\widetilde{H}_*(\widetilde{S(\Gamma)}_t) \cong \bigoplus_{v \notin \widetilde{S(\Gamma)}_t} \widetilde{H}_*(F(\Gamma)).$$

The following is now immediate:

Corollary 3.22. *The group K_Γ is of type $FH_{n+1}(R)$ if and only if $F(\Gamma)$ is homologically n–connected over R and is of type $FH(R)$ if and only if $F(\Gamma)$ is acyclic over R.*

Right-angled Artin groups themselves are $FH(R)$ over every ring. So, if $F(\Gamma)$ is not acyclic over some ring R then K_Γ cannot be isomorphic to a right-angled Artin group.

It is also true that K_Γ is finitely presented if and only if $F(\Gamma)$ is simply connected, but we will not give the full proof here. Instead, we will direct the reader to [4].

4. Surface subgroups of right-angled Artin groups

The main goal of this section is to understand how long cycles in the underlying graph Γ give rise to closed surface subgroups of right-angled Artin groups. We will present two perspectives on this fact – one due to Droms, Servatius and Servatius which constructs surface subgroups directly inside the right-angled Artin group, and one due to Davis and Januszkiewicz, which takes a detour through right-angled Coxeter groups.

4.1. **Right-angled Artin groups on cycles which do not contain surface groups.** First, we observe that if $n \leq 4$ and C_n denotes the cycle of length n, then $A(C_n)$ does not contain a closed hyperbolic surface group. In the case where $n = 2, 3$, we have that $A(C_n)$ is abelian so that we are done. If $n = 4$, we have that $A(C_4) \cong F_2 \times F_2$. The strategy of proof for the following proposition is useful in the analysis of right-angled Artin groups on joins:

Proposition 4.1. *The group $F_2 \times F_2$ does not contain a closed hyperbolic surface group.*

Proof. Let $G = \pi_1(S)$, where S is closed and of genus at least two, and let $\phi : G \to F_2 \times F_2$ be an injective homomorphism. Write p_1 and p_2 for the projections of $F_2 \times F_2$ onto the two direct factors. Since G is not free, we have that both $p_1 \circ \phi$ and $p_2 \circ \phi$ are not injective. Write K_1 and K_2 for the kernels of these maps, both of which are nontrivial.

Observe that since ϕ is injective, we have that $K_1 \cap K_2 = \{1\}$. Furthermore, K_1 and K_2 centralize each other. It follows that $K_1 K_2 < F_2 \times F_2$ is isomorphic to $K_1 \times K_2$. It follows that

$$\mathbb{Z}^2 < K_1 \times K_2 < G,$$

a contradiction. □

4.2. **The Droms–Servatius–Servatius construction.** Let $w \in A(\Gamma)$ be a word in the vertices of Γ which represents the identity in $A(\Gamma)$:

Proposition 4.2. *The word w can be reduced to the identity after performing a finite sequence of the following two types of moves:*

1. *Canceling subwords of the form $x \cdot x^{-1}$.*
2. *Switching two consecutive vertices $v_1 \cdot v_2 \mapsto v_2 \cdot v_1$ when $(v_1, v_2) \in E(\Gamma)$.*

Proof. Exercise, or consult [41]. □

Let $S(\Gamma)$ be the Salvetti complex of Γ, and let $X(\Gamma)$ denote its 2–skeleton. We have that $X(\Gamma)$ is the Cayley 2–complex associated to the standard presentation of $A(\Gamma)$. Let $\widetilde{X} \to X(\Gamma)$ be any covering space (with the pulled–back cell structure) and let $\gamma \subset \widetilde{X}$ be a loop in the 1–skeleton of $\widetilde{X}$ which is homotopic to a point. Then γ is homotopic to a point via a finite sequence of paths

$$\gamma = \gamma_0 \to \gamma_1 \to \cdots \to \gamma_n,$$

where γ_n is a point, and where each of the transitions $\gamma_i \to \gamma_{i+1}$ is given by eliminating a backtracking path (a move of type (1) in Proposition 4.2), or by "sliding two edges of γ_i across a square" in $\widetilde{X}$ (a move of type (2) in Proposition 4.2). An immediate consequence of this discussion is the following observation, which gives us a criterion for checking the π_1–injectivity of the inclusion of a subcomplex of $\widetilde{X}$:

Proposition 4.3. *Let $\widetilde{X} \to X(\Gamma)$ be a covering space and let $Z \subset \widetilde{X}$ be a subcomplex. Suppose that if $F \subset \widetilde{X}$ is a face which contains at least two incident edges of Z, then F is also a face of Z. Then the inclusion of Z into $\widetilde{X}$ induces an injection of fundamental groups.*

If Z has the property described in the previous proposition, we will say that Z is a *full* subcomplex of $\widetilde{X}$.

The following result is now remarkably straightforward, though incredibly slick:

Theorem 4.4 (See [20])**.** *Let $n \geq 5$ and let C_n be a cycle of length n. Then $A(C_n)$ contains a closed hyperbolic surface group.*

Proof. First, recall that the Salvetti complex of C_n can be built by taking a union of distinguished subcubes of the unit cube in $\mathbb{R}^n$ and then by taking the quotient by the $\mathbb{Z}^n$–action on $\mathbb{R}^n$, so that $X(C_n) \subset S(C_n) \subset \mathbb{R}^n/\mathbb{Z}^n$. We now take the cover of $X(C_n)$ induced by taking the universal cover of $\mathbb{R}^n/\mathbb{Z}^n$, and we call it $\widetilde{X}$. Clearly we have an isomorphism

$$\pi_1(\widetilde{X}) \cong [A(C_n), A(C_n)].$$

Let $I^n \subset \mathbb{R}^n$ be the unit cube. We now consider $S = I^n \cap \widetilde{X}$. Observe that S is a subcomplex of the 2–skeleton of I^n. Note that each edge of I^n corresponds to the (deck transformation of $\mathbb{R}^n$) induced by a vertex of C_n, by construction. Each edge of I^n is incident to $n-1$ other edges and hence to $n-1$ faces. Observe that in C_n, each vertex is adjacent to exactly two other vertices. It follows that in S, each edge is incident to exactly two other edges and is hence incident to exactly two faces. It follows that S is a full and connected subcomplex of I^n. So, S is a connected 2–complex in which each edge is incident to exactly two faces and $\pi_1(S)$ injects into $\pi_1(\widetilde{X})$, so that S is a closed surface whose fundamental group is a subgroup of $A(C_n)$. Since it is a subcomplex of I^n, the complex S is two–sided.

One easily computes the Euler characteristic of S to be $(4-n)2^{n-2}$, so that the genus of S is $1+(n-4)2^{n-3}$. Thus, the genus of S is greater than two whenever $n \geq 5$. □

The previous result can be used to prove a curious fact concerning the structure of commutator subgroups of right-angled Artin groups. A graph is called *triangulated* if it contains no induced copy of C_n for $n \geq 4$.

Lemma 4.5. *Let Γ be a finite, connected, triangulated graph with at least two vertices. Then Γ is a union of two proper induced subgraphs X and Y such that $X \cap Y \neq X, Y$ is complete.*

The proof of the previous lemma is left as a (not entirely trivial) exercise for the reader.

Theorem 4.6 (See [20])**.** *Let Γ be a finite simplicial graph. The derived subgroup $[A(\Gamma), A(\Gamma)]$ is free if and only if Γ is a triangulated graph.*

Proof. If Γ is not triangulated then we have an induced copy of C_n inside of Γ for some $n \geq 4$. If $n > 4$ then we have seen that $[A(\Gamma), A(\Gamma)]$ contains a closed hyperbolic surface group, so that $[A(\Gamma), A(\Gamma)]$ is not free. The proof also shows that $[A(C_4), A(C_4)]$ contains a copy of $\mathbb{Z}^2$, so that in this case we also have that $[A(\Gamma), A(\Gamma)]$ is not free. That $[A(C_4), A(C_4)]$ contains a copy of $\mathbb{Z}^2$ can also be seen directly: if we write

$$F_2 \times F_2 \cong \langle a, b\rangle \times \langle c, d\rangle,$$

then the elements $[a,b] \times \{1\}$ and $\{1\} \times [c,d]$ are clearly in the derived subgroup and together generate a copy of $\mathbb{Z}^2$.

The proof of the converse is by induction on the number of vertices of Γ, the base case of one vertex being trivial. If Γ is disconnected, then any subgroup of $A(\Gamma)$ is a free product of a free group and conjugates of subgroups of the right-angled Artin groups on the components, by the Kurosh subgroup theorem. If $H < [A(\Gamma), A(\Gamma)]$ and H_i is a free factor of H which is conjugate into one of the component right-angled Artin groups, say $A(\Gamma_i)$, then $H_i < [A(\Gamma_i), A(\Gamma_i)]$.

If Γ_i is a component of Γ, the fact that Γ is triangulated implies that $\Gamma_i = X \cup Y$ with $X \cap Y \neq X, Y$ a complete subgraph. It follows that

$$A(\Gamma_i) \cong A(X) *_{A(X\cap Y)} A(Y),$$

and that X and Y have fewer vertices than Γ_i. Observe that for each $g \in A(\Gamma)$, we have that

$$[A(\Gamma_i), A(\Gamma_i)] \cap A(X \cap Y)^g = \{1\},$$

since $A(X \cap Y)$ injects into the abelianization $A(\Gamma)^{ab}$. It follows that $[A(\Gamma_i), A(\Gamma_i)]$ acts on the Bass–Serre tree associated to the amalgamated splitting of $A(\Gamma_i)$ above, with trivial edge stabilizers. It follows that $[A(\Gamma_i), A(\Gamma_i)]$ is a free product of subgroups conjugate into $[A(X), A(X)]$ or $[A(Y), A(Y)]$, possibly with a free group factor. By induction, it follows that $[A(\Gamma_i), A(\Gamma_i)]$ is free. □

4.3. **Surface subgroups via right-angled Coxeter groups.** In this subsection we will discuss a result of M. Davis and T. Januszkiewicz from [16]. Their result, as the title suggests, is that each right-angled Artin group is commensurable with a right-angled Coxeter group. In particular, they conclude that right-angled Artin groups are linear over $\mathbb{Z}$ and hence have all sorts of nice properties.

We will require a more precise statement of their theorem:

Theorem 4.7. *Let Γ be a finite simplicial graph. There exists a finite simplicial graph Γ' which contains Γ as an induced subgraph, such that the group $A(\Gamma)$ is commensurable with the right-angled Coxeter group $C(\Gamma')$.*

The reason we want the more precise statement is the following:

Corollary 4.8. *For any $n \geq 5$, the right-angled Artin group $A(C_n)$ contains a closed hyperbolic surface group.*

Proof. By Theorem 4.7, it suffices to show that the right-angled Coxeter group $C(C_n)$ contains a closed hyperbolic surface group. This follows from the fact that $C(C_5)$ is the group of reflections in the sides of a right-angled pentagon in $\mathbb{H}^2$, and that a closed genus two surface can be tiled with eight right-angled pentagons. Furthermore, the groups $C(C_n)$ are all commensurable with each other for $n \geq 5$ (this will be discussed more in the sequel). The claim follows. □

We will now fix Γ and build two right-angled Coxeter groups on graphs X and Y which are closely related to Γ. Write $V(\Gamma)$ for the vertex set of Γ. We define $V(X) = V(\Gamma) \times \{0, 1\}$. We let the subgraph of X spanned be $V(\Gamma) \times \{1\}$ be Γ, and the subgraph of X spanned by $V(\Gamma) \times \{0\}$ be complete. We then connect $(v_1, 0)$ and $(v_2, 1)$ in X if and only if $v_1 \neq v_2$.

We define $V(Y) = V(\Gamma) \times \{-1, 1\}$. The subgraphs spanned by $V(\Gamma) \times \{1\}$ and $V(\Gamma) \times \{-1\}$ in Y are both defined to be Γ. We connect $(v_1, -1)$ with $(v_2, -1)$ in Y whenever v_1 and v_2 are adjacent in Γ. Let $\{r_i\}_{i \in V(\Gamma)}$ be the usual generating set for the direct sum $(\mathbb{Z}/2\mathbb{Z})^{V(\Gamma)}$. Following Davis and Januszkiewicz, we define two homomorphisms

$$\phi, \theta : C(X) \to (\mathbb{Z}/2\mathbb{Z})^{V(\Gamma)}$$

by

$$\phi : (v_i, 0) \mapsto r_i$$

and

$$\phi : (v_i, 1) \mapsto 0,$$

and by

$$\theta : (v_i, 0), (v_i, 1) \mapsto r_i.$$

We now define homomorphisms $\alpha : C(Y) \to \ker \phi < C(X)$ and $\beta : A(\Gamma) \to \ker \theta < C(X)$ by

$$\alpha : (v_i, 1) \mapsto (v_i, 1)$$

and

$$\alpha : (v_i, -1) \mapsto (v_i, 1)^{(v_i, 0)},$$

and by

$$\beta : v_i \mapsto (v_i, 0)(v_i, 1).$$

The main result of Davis–Januszkiewicz can be stated as follows:

Theorem 4.9. *The homomorphisms α and β above are isomorphisms.*

One easily checks that the graph X contains the graph Γ as an induced subgraph, so that this result is sufficient in order to establish the existence of closed hyperbolic surface groups inside of $A(C_n)$ for $n \geq 5$.

The idea behind Theorem 4.9 is to produce explicit cube complexes on which $A(\Gamma)$, $C(X)$ and $C(y)$ act, and to analyze to what degree these complexes and the associated actions are geometrically similar. Before giving a proof of the result, we will discuss some constructions.

Let Z be a topological space. An *mirror structure* on Z over an indexing set V is a family $\mathcal{M} = (Z_i)_{i \in V}$ of closed subspaces of Z indexed by the set V. We let Γ be a graph on V and $C(\Gamma)$ the associated right-angled Coxeter group. If a mirror sturcture on Z is given, we can define an equivalence relation on $C(\Gamma) \times Z$ by $(w, z) \sim (w', z')$ whenever $z = z'$ and $w^{-1}w'$ belongs to the subgroup of W generated by the vertices $\{s_i\}$ of Γ for which $z \in Z_i$. We will write U for the quotient of $C(\Gamma) \times Z$ by this equivalence relation.

Observe that $C(\Gamma)$ acts on U and U has the following universal property: if Y is any $C(\Gamma)$–space and if $f : Z \to Y$ is any map such that $f(Z_i)$ is contained in the fixed set of s_i for each i, then f extends to a $C(\Gamma)$–equivariant map $U \to Y$ by the assignment $(w, z) \mapsto w \cdot f(z)$.

Let $F(\Gamma)$ be the flag complex of Γ as before, and let $P(\Gamma)$ be the *poset* associated to Γ. The latter is just a partially ordered set on the complete subgraphs of Γ, partially ordered by inclusion.

We will be concerned with the posets $P(X)$, $P(Y)$ and P associated to the graphs X, Y and Γ above, respectively. The cubical complexes $K(X)$ and $K(Y)$ are defined by

$$K = [0, 1]^V \cap \bigcup_{J \in P} \mathbb{R}^J$$

and

$$K(Y) = [-1, 1]^V \cap \bigcup_{J \in P} \mathbb{R}^J.$$

Here, the notation $\mathbb{R}^J$ means the subspace of $\mathbb{R}^V$ given by setting the coordinates indexed by $V \setminus J$ to be zero.

Now for a subset $J \subset V$, write

$$K_J = K \cap \{x_j = 1\}_{j \in J}.$$

Let J' be a subset of $I \times \{-1, 1\}$ defined by the graph of a function $\epsilon : J \to \{-1, 1\}$, and let

$$K(Y)_{J'} = K(Y) \cap \{x_j = \epsilon_j\}_{(j, \epsilon_j) \in J'}.$$

Finally, let $J'' = J^0 \times 0 \cup J^1 \times 1$ be a subset of $I \times \{0, 1\}$ and let

$$K(X)_{J''} = K \cap \mathbb{R}^{J^0} \cap \{x_j = 1\}_{j \in J^1}.$$

We have that $K(Y)_{J'}$ is nonempty if and only if $J' \in P(Y)$ and $K(X)_{J''}$ is nonempty if and only if $J^0 \cap J^1 = \emptyset$ and $J^1 \in P$, which is to say if and only if $J'' \in P(X)$.

For each $v \in V$, we let $K_v = K_{\{v\}}$, which defines a mirror structure $\mathcal{M}$ on K. We similarly obtain mirror structures on $K(X)$ and on $K(Y)$. We will write U, $U(X)$ and $U(Y)$ for the associated natural cube complexes on which $C(\Gamma)$, $C(X)$ and $C(Y)$ act with compact quotient.

Lemma 4.10. *The complexes U, $U(X)$ and $U(Y)$ are contractible.*

Proof. See [16]. □

Lemma 4.11. *The map α in the statement of Theorem 4.9 is an isomorphism. Furthermore, there is an α–equivariant homeomorphism $U(Y) \to U(X)$.*

Proof. The basic idea is to decompose $K(Y)$ as a union of copies of $K(X)$ in a way which respects the $\mathbb{Z}/2\mathbb{Z}^V$ action of $[-1, 1]^V$, given by reflection about the planes $x_i = 0$. Since this result is peripheral to our purposes, we refer the reader to its proof in [16]. □

Lemma 4.12. *The map β in the statement of Theorem 4.9 is an isomorphism. The cubical complex $U(X)$ can be identified with the universal cover of the Salvetti complex $S(\Gamma)$.*

Proof. We have already seen a canonical way to build $S(\Gamma)$ as a subcomplex of the torus $(S^1)^V$. For our purposes, we build $S(\Gamma)$ in the same way as a quotient of the complex $[0,2]^V$ with opposite faces identified, which we denote by T^V.

As before, we let $\{r_i\}$ be the standard generating set for $\mathbb{Z}/2\mathbb{Z}^V$. We let r_i act on T^V by reflection across the hyperplane $x_i = 1$. Observe that this defines an action of $\mathbb{Z}/2\mathbb{Z}^V$ on T^V, and $[0,1]^V$ is the fundamental chamber of this action. Furthermore, $S(\Gamma)$ is stable under this action.

The complex $K = K(X)$ above is a fundamental chamber for the $\mathbb{Z}/2\mathbb{Z}^V$ action on $S(\Gamma)$. We will write $K(A) = K$ and $\mathcal{M}(A)$ for the mirror structure given by this action. Write $U(A)$ for the associated cube complex.

There is an identification of $U(A)$ with $S(\Gamma)$, more or less by definition. The inclusion $K(X) \to U(A)$ and the map $\theta : C(X) \to \mathbb{Z}/2\mathbb{Z}^V$ gives a θ–equivariant map $U(X) \to U(A)$ which is a covering map. So, we have homeomorphisms

$$S(\Gamma) \cong U(A) \cong U(X)/\ker\theta.$$

We get the conclusion of the lemma by lifting to universal covers. □

4.4. **One–ended subgroups of right-angled Artin groups.** Now that we know that surface groups embed generally embed in right-angled Artin groups, we would like to know what their images may look like. We will give a discussion here which applies to right-angled Artin groups but which arises from a deep question about mapping class groups. We will return to this discussion in the sequel after we have developed more mapping class group–like machinery for right-angled Artin groups.

Closed surface groups are examples of *one–ended* groups, since their Cayley graphs have exactly one end. Let G be a finitely presented, one–ended group, and suppose $G < A(\Gamma)$ for some graph Γ. The following result is due to C. Leininger:

Theorem 4.13 (See [35])**.** *There exists an element $1 \neq g \in G$ such that $g \in \mathbb{Z}^2 < A(\Gamma)$.*

In the case where $\Gamma = C_n$ for some $n \geq 5$, the Centralizer Theorem implies that, up to conjugacy, there is an element $1 \neq g \in G$ which is supported in the star of some vertex of C_n. Observe that the star of a vertex v generates a subgroup of $A(C_n)$ isomorphic to $F_2 \times \mathbb{Z}$, so that $g = w \cdot z$ where $w \in F_2$ and z is contained in the central copy of $\mathbb{Z}$.

The Droms–Servatius–Servatius result exhibits a closed hyperbolic surface group contained in $[A(C_n), A(C_n)]$. Observe that if $g = w \cdot z$ and is also contained in the commutator subgroup, then z is trivial and w is contained in the commutator subgroup of F_2. Thus, Leininger's result shows that the Droms–Servatius–Servatius surface subgroups of $A(C_n)$ contain nontrivial elements of a very restricted form.

If $v \in V(\Gamma)$ is a vertex, which as usual we identify with a generator of $A(\Gamma)$ and as a 1–cell of $S(\Gamma)$, we can define a *dual hypersurface* to v, which we denote by T_v. In $\mathbb{R}^V$, we let A_v be the subspace defined by $x_v = 1/2$. Recall that $S(\Gamma)$ was constructed as a quotient of a subcomplex of $[0,1]^V$, and we define T_v to be the image of $A_v \cap [0,1]^V$ inside of $S(\Gamma)$.

Observe that T_v is nonempty an it locally embeds as a codimension one subtorus of standard torus and has a collar neighborhood $N(T_v)$ which is homeomorphic to $T_v \times [-1,1]$. In general, T_v is not a torus but a union of tori. Observe that any element of $\pi_1(T_v) < A(\Gamma)$ commutes with v. In particular, whenever T_v has positive dimension, each element of $\pi_1(T_v)$ has noncyclic stabilizer.

Proof of Theorem 4.13. We follow Leininger's proof. Let X be a presentation 2–complex for G, so that $G = \pi_1(X)$. Let $f : X \to S(\Gamma)$ be a map which, after fixing basepoints, realizes the inclusion of G into $A(\Gamma)$.

Label the vertices of Γ by $\{v_1, \ldots, v_n\}$, and we will write $T_i = T_{v_i}$. We may assume that f is transversal to T_1 after possibly applying a homotopy. We may assume that $f^{-1}(T_1)$ is nonempty, for otherwise we could express G as a subgroup of the right-angled Artin subgroup of $A(\Gamma)$ generated by $\{v_2, \ldots, v_n\}$. Since there are no one–ended subgroups of $\mathbb{Z}$, we may apply induction on the number of vertices of Γ to assume $f^{-1}(T_1)$ is nonempty, and that the number of components of $f^{-1}(T_1)$ is minimal within the homotopy class of f.

Observe that $f^{-1}(T_1)$ is a subgraph of X which is locally separating (since each T_i separates the neighborhood $N(T_i)$). We will write Y for a component of $f^{-1}(T_1)$. If Y contains a homotopically nontrivial loop γ, then the fact that induced map f_* of fundamental groups between G and $A(\Gamma)$ is injective implies that $f_*([\gamma]) \neq 1$. In particular, T_1 must have positive dimension, and the centralizer of $f_*([\gamma])$ is contains a copy of $\mathbb{Z}^2$.

So, we may suppose that every loop in Y is homotopically trivial. We may therefore lift Y to the universal cover $\widetilde{X}$ of X. Observe that any lift of Y is a compact subgraph of $\widetilde{X}$. We will choose a lift of Y and abuse notation by calling it Y as well. The fact that Y has a bi–collar in X implies that Y separates $\widetilde{X}$ as well. Indeed, suppose not. Then we can cut open $\widetilde{X}$ along a lift of Y and glue together infinitely many copies of the resulting space together, which furnishes us with a connected, infinite cyclic cover of $\widetilde{X}$. Since $\widetilde{X}$ is simply connected, this is a contradiction.

So, any lift of Y cuts $\widetilde{X}$ into two components. Note that at most one of these components can be noncompact, since otherwise the compactness of Y implies that $\widetilde{X}$ has at least two ends. Let $Z \subset \widetilde{X}$ be the bounded component. We have that $N(Y) \cup Z$ is compact, since $N(Y)$ is compact and since the closure of Z is $Y \cup Z$.

Now we claim that the covering map $\widetilde{X} \to X$ is a homeomorphism when restricted to $N(Y) \cup Z$. It suffices to show that any nontrivial deck transformation of $\widetilde{X}$ over X takes $N(Y) \cup Z$ off of itself. Observe first that any nontrivial covering transformation of $\widetilde{X}$ takes Y to a different (and hence disjoint) lift of Y. Therefore, if γ is a nontrivial deck transformation, we may assume $\gamma(N(Y)) \cap N(Y) = \emptyset$. So, suppose that $\gamma(Z) \cap Z \neq \emptyset$. It follows that Z is contained in the bounded component of $X \setminus \gamma(Y)$. Replacing γ by γ^{-1}, we may assume $\gamma^k(Z) \subset Z$ for all k, which is impossible by the proper discontinuity of the deck group action and the compactness of Z. This establishes the claim.

It follows that the map f can be homotoped in a way to decrease the number of components of $f^{-1}(T_1)$. Indeed, we can homotope f by a homotopy supported on $Z \cup N(Y)$ to miss T_1 completely, since $Z \cup N(Y)$ lands in a contractible subset of $S(\Gamma)$. We have thus contradicted the minimality assumption on the number of components of $f^{-1}(T_1)$, so that we Y must contain a homotopically nontrivial loop. □

Leininger's argument also gives a fresh perspective on two–generated subgroups of right-angled Artin groups, which as we have seen are either abelian and free. Aside from two technical issues, his argument gives a new, efficient proof of that result.

Lemma 4.14 (See [42]). *Let G be a two–generated subgroup of $A(\Gamma)$. If G has infinitely many ends then G is free.*

Proof. By Stallings' work (see [42]), we have that a torsion–free group with more than one end is either cyclic or admits a decomposition as a nontrivial free product. We have that $G \cong A * B$, where both A and B are not trivial. Observe that if A and B are nontrivial then they both admit cyclic quotients, since right-angled Artin groups are residually p for each prime p. It follows that G admits a surjection to F_2. The Hopficity of free groups implies that $G \cong F_2$. □

Theorem 4.15. *Let $G < A(\Gamma)$ be finitely presented and two–generated, and suppose that the diameter of Γ is at least five. Then G is either abelian or free.*

Proof. Assume G is nonabelian. Evidently we may assume that G is one–ended, so that all the machinery in the proof of Theorem 4.13 applies. As before, we write X for a finite Cayley 2–complex for G and

$$f : X \to S(\Gamma)$$

a transverse map realizing the inclusion of G into $A(\Gamma)$. By induction, we may assume that Γ is connected and does not split as a join. Write $T_1, \ldots, T_k$ for the dual hypersurfaces in $S(\Gamma)$. Again by induction we may assume that for each i, the preimage $f^{-1}(T_i)$ is nonempty and that each component of $f^{-1}(T_i)$ contains a homotopically nontrivial loop in X (up to free homotopy).

Observe that for each i, the support of $\pi_1(T_i)$ is contained in the link of the vertex v_i. Since the diameter of Γ is at least five, there are two vertices, say v_1 and v_2, such that $st(v_1) \cap st(v_2) = \emptyset$. The projection

$$\phi : A(\Gamma) \to \langle st(v_1) \cup st(v_2)\rangle$$

furnishes a map from $A(\Gamma)$ to a smaller right-angled Artin group, and furthermore $\phi(G)$ contains a nonabelian free group and is hence a two–generated, nonabelian subgroup of a smaller right-angled Artin group. The result follows by induction on the number of vertices of Γ. □

5. Right-angled Artin subgroups of mapping class groups

We will now give a mostly complete description of right-angled Artin subgroups of mapping class groups, and we will develop some mapping class group tools which can be used to study right-angled Artin groups for their own sake. The main references for this section will be [32] and [31]. Unless otherwise noted, all statements and proofs can be found in those two references.

5.1. **Generalities on mapping class groups.**

Let S be an orientable surface of genus g and n punctures, so that

$$0 < 2g - 2 + n < \infty.$$

The mapping class group of S is defined by

$$\mathrm{Mod}(S) = \mathrm{Homeo}^+(S)/\sim,$$

where the relation we place on the orientation–preserving homeomorphisms of S is isotopy, which is to say homotopy through homeomorphisms. Mapping class groups are finitely presented groups with various interesting properties. For a very thorough introduction, the reader is directed to [21].

Let $c \subset S$ be an *essential, nonperipheral simple closed curve*, which we abbreviate *scc*. By this we mean that c is a smoothly embedded copy of $S^1 \subset S$ such that the free homotopy class of c is nontrivial, and such that c is not isotopic to a puncture. We will generally abuse terminology and identify a scc with its free homotopy class. Since mapping classes are homeomorphisms of S up to isotopy (or homotopy – replacing equivalence of homeomorphisms by isotopy with homotopy results in the same mapping class group), we get an action of $\mathrm{Mod}(S)$ on the set of sccs. The sccs on S are useful for describing a trichotomy among mapping classes. The following is a well–known definition–theorem:

Theorem 5.1 (See [22]). *Let* $1 \neq \psi \in \mathrm{Mod}(S)$.

(1) *The mapping class ψ has finite order if and only if for every scc c there is a $k \neq 0$ such that $\psi^k(c) = c$. Equivalently, there is a $k \neq 0$ such that for each scc c we have $\psi^k(c) = c$. If ψ has finite order then there is a hyperbolic metric on S for which a representative of ψ acts by an isometry.*

(2) *The mapping class ψ is infinite–order reducible if it does not have finite order and if there is a nonempty, finite union of disjoint sccs C (also called a multicurve) such that $\psi(C) = C$. The minimal nonempty such multicurve is called a canonical reduction system for ψ.*

(3) *The mapping class* ψ *is pseudo-Anosov if it is neither finite order nor reducible. Equivalently, there is a pair of singular, measured foliations* $\mathcal{F}^{\pm}$ *with dense leaves which are preserved by* ψ *and are stretched/contracted by a positive, real scalar* $\lambda^{\pm 1}$*, where* $\lambda > 1$.

An example of a finite order mapping class can be produced by arranging a finite number of handles around a central torus. A finite order homeomorphism of the surface is given by rotating the surface one "click".

A common example of a reducible mapping class is given by a *Dehn twist* about a scc. The Dehn twist about an scc c is defined by cutting S open along c and then regluing S after rotating one of the boundaries by a full twist.

Pseudo-Anosov homeomorphisms, despite being the most common (in some sense) mapping classes, are the most difficult mapping classes to exhibit. A construction due to Thurston goes as follows: let c, c' be two sccs which fill S, which is to say any other scc c'' is forced to intersect at least one of c and c'. Then for all sufficiently large N, the elements $T_c^N T_{c'}^N$ are all pseudo-Anosov.

A nontrivial mapping class ψ is called *pure* if it is pseudo-Anosov or if it is reducible and if it preserves the components of its canonical reduction system C, it preserves the components of $S \setminus C$, and if the restriction of ψ to any component of $S \setminus C$ is either pseudo-Anosov or the identity. If ψ is a mapping class such that ψ is pseudo-Anosov when restricted to a connected subsurface $S_0 \subset S$ and restricts to the identity on $S \setminus S_0$, we will say that ψ is *pseudo-Anosov on a subsurface*. By convention, we consider pseudo-Anosov mapping classes to be pseudo-Anosov on a subsurface.

Reducible mapping classes can give rise to abelian subgroups of mapping class groups. A result of Birman, Lubotzky and McCarthy gives control over the size of abelian subgroups of $\mathrm{Mod}(S)$;

Theorem 5.2 ([6])**.** *Every virtually solvable subgroup of* $\mathrm{Mod}(S)$ *is virtually abelian. If* $A < \mathrm{Mod}(S)$ *is torsion–free and abelian then* $\operatorname{rk} A$ *is bounded by the maximal number of components in a multicurve on* S.

5.2. **Curve complexes.** It is useful to organize all the sccs on a surface S into a topologically natural simplicial complex. Such a complex would then admit an action of $\mathrm{Mod}(S)$ and aid in its study. Let $S = S_{g,n}$, and suppose $S \notin \{S_{0,3}, S_{0,4}, S_{1,1}\}$. The *curve graph* $\mathcal{C}(S)$ of S is defined as follows: the vertices of $\mathcal{C}(S)$ are sccs, and two sccs are adjacent if they can be realized disjointly in S. The *curve complex* of S is the flag complex of the curve graph. In the sequel, we will only be using the curve graph, so we will not define notation for the curve complex of S.

Since the mapping class group acts on free homotopy classes of curves on S and since its elements are homotopy classes of homeomorphisms, the adjacency relation of $\mathcal{C}(S)$ is preserved by $\mathrm{Mod}(S)$. It follows that $\mathrm{Mod}(S)$ acts on $\mathcal{C}(S)$ by graph automorphisms. Evidently there is only one nonseparating curve on S up to homeomorphism and only finitely many separating curves on S up to homeomorphism. It follows that the quotient of $\mathcal{C}(S)$ by $\mathrm{Mod}(S)$ is finite.

The curve graph is trivial for $S_{0,3}$ since there are no sccs, and the curve graph can be defined for $S_{0,4}$ and $S_{1,1}$. Since there exists no pair of disjoint sccs on either of these surfaces, the adjacency relation is defined to be by minimal intersection. That is to say, two sccs on these surfaces are adjacent if they admit representatives which meet a minimal number of times. In the case of $S_{0,4}$ this means two intersections, and in the case of $S_{1,1}$ this means one intersection.

The curve graph is an extraordinarily complicated object, though it does have many useful finiteness properties. One such finiteness property is that the curve complex $\mathcal{C}(S)$ has a finite chromatic number. Recall that a graph X has chromatic number at most χ if there exists a map $\mathit{color} : X \to S$, where S is a set of cardinality χ, with the property that *color* does not assign the same value to adjacent vertices. Such a map is called a *coloring* of X. The *chromatic number* $\chi(X)$ of X is defined to be the smallest cardinality of such a set S.

The following result is elementary, though it is very useful and was only proved first by Bestvina, Bromberg and Fujiwara in [5]:

Proposition 5.3. *For any surface* S*, the curve graph* $\mathcal{C}(S)$ *has a finite chromatic number.*

Proof. We will assume $S \notin \{S_{0,3}, S_{0,4}, S_{1,1}\}$, leaving the remaining cases to the reader. Given S, we consider the (finite) collection $\{S_i\}$ of all connected degree two covers of S. Let

$$F = \bigoplus_i (H_1(S_i, \mathbb{Z}/2\mathbb{Z})^{\oplus 2} \oplus \{0\}),$$

where we consider $H_1(S_i, \mathbb{Z}/2\mathbb{Z})^{\oplus 2}$ as unordered pairs of homology classes. Evidently F is a finite set.

We let *color* : $\mathcal{C}(S) \to F$ be defined coordinate–wise. If γ does not lift to S_i, we defined *color*$(\gamma) = 0$ in that coordinate. Otherwise, we send γ to the unordered pair of modulo two homology classes of its two lifts to S_i. We claim that *color* is a coloring of $\mathcal{C}(S)$. There are only finitely many cases that need to be checked, since there are only finitely many configurations of pairs (γ_1, γ_2) of disjoint sccs on S, up to homeomorphism. Namely:

(1) Both γ_1 and γ_2 are separating.
(2) The curve γ_1 is nonseparating and γ_2 is separating.
(3) Both γ_1 and γ_2 are nonseparating but their union is separating.
(4) Both γ_1 and γ_2 are nonseparating and their union is nonseparating.

To prove the proposition, it suffices to find a degree two cover of S where modulo two homology distinguishes between the two curves, so that *color* assigns different colors to γ_1 and γ_2. For example, in case (1) it is possible to find a degree two cover of S where both lifts of γ_1 are separating but neither lift of γ_2 is separating (so that the modulo two homology classes of the lifts are both nonzero). We leave the details of the verification to the reader. □

5.3. **Clique graphs.** Let X be any graph. We will write X_k for the *clique graph* of X. The vertices $\{v_K\}$ of X_k are in bijection with the complete subgraphs $\{K\}$ of X with at least one vertex, and two distinct vertices v_{K_1} and v_{K_2} are adjacent if and only if K_1 and K_2 span a complete subgraph of X.

Proposition 5.4. *Let X be a graph with a finite chromatic number. Then the chromatic number of X_k is also finite.*

Proof. Exercise. □

Observe that X_k contains a natural copy of X. If X has no triangles then X_k is isomorphic to X together with a triangle hanging off every edge of X like a fin.

5.4. **Projective measured laminations.** The sccs on a surface S can be organized into another object, the *projective measured laminations* on S, denoted $\mathbb{P}\mathcal{ML}(S)$. This object is connected with the curve graph of S, though not in a straightforward way. Fixing a hyperbolic metric on S, we think of elements of $\mathbb{P}\mathcal{ML}(S)$ as unions of geodesics equipped with a projective class of recurrent transverse measures. The space $\mathbb{P}\mathcal{ML}(S)$ has a natural piecewise–integral–linear structure, and the mapping class group acts on $\mathbb{P}\mathcal{ML}(S)$ by piecewise–linear homeomorphisms. We will not digress into a general discussion of $\mathbb{P}\mathcal{ML}(S)$. The reader will find all the necessary facts in Penner and Harer's volume [36].

A pure mapping class has a collection of invariant measured foliations. If ψ is a Dehn twist about a scc then we define the twisting curve to be the *limiting lamination* of ψ. If ψ is pseudo-Anosov on a (possibly proper) subsurface $S_0 \subset S$, then we call the stable and unstable laminations $\lambda^{\pm}$ of ψ on S_0 the limiting laminations of ψ.

5.5. **Ping–pong and free subgroups.** The general theory of projective measured laminations implies that if members of a collection of mapping classes $\{f_1, \ldots, f_k\}$ each have connected supports and pairwise intersecting limiting laminations then they will stably generate a free group:

Theorem 5.5. *Let $\{f_1, \ldots, f_k\} \subset \mathrm{Mod}(S)$ be a collection of Dehn twists about simple closed curves and pseudo-Anosov mapping classes on subsurfaces with limiting laminations $\{\mathcal{L}_1^{\pm}, \ldots, \mathcal{L}_k^{\pm}\}$. Suppose that for each $i \neq j$, we have $\mathcal{L}_i^{\pm} \neq \mathcal{L}_j^{\pm}$ and that $\mathcal{L}_i^{\pm} \cap \mathcal{L}_j^{\pm} \neq \emptyset$. Then there exists an N such that for all $n \geq N$, we have*

$$\langle f_1^n, \ldots, f_k^n \rangle \cong F_k.$$

Proof. Let γ be a simple closed curve on S, viewed as a projective measured lamination, which intersects each of the limiting laminations of the mapping classes in question. Let $\{U_1, \ldots, U_k\}$ be neighborhoods of $\{\mathcal{L}_1^{\pm}, \ldots, \mathcal{L}_k^{\pm}\}$ in $\mathbb{P}\mathcal{ML}(S)$ whose union does not contain γ. Observe that for each sufficiently large N and for each i, we have $f_i^{\pm N}(\gamma) \in U_i$. Furthermore, for each $i \neq j$ and each sufficiently large N, we have $f_i^{\pm N}(U_j) \subset U_i$. The conclusion follows by an easy application of the ping–pong lemma. □

One can weaken the hypotheses on $\{f_1, \ldots, f_k\}$ slightly. One can assume that the supports of $\{f_1, \ldots, f_k\}$ are disconnected, provided that each $i \neq j$, each leaf of $\mathcal{L}_i^{\pm}$ intersects each leaf of $\mathcal{L}_j^{\pm}$. We leave it as an exercise for the reader to check that Theorem 5.5 is false if one merely requires $\mathcal{L}_i^{\pm} \cap \mathcal{L}_j^{\pm} \neq \emptyset$.

Observe that if a collection $\{f_1, \ldots, f_k\} \subset \mathrm{Mod}(S)$ generate a subgroup isomorphic to F_k then the same is true if we replace these mapping classes by any nonzero power. In view of this observation and the exercise in the previous paragraph, we see that Theorem 5.5 gives the most general characterization of the structure of free subgroups of mapping class groups.

5.6. Right-angled Artin subgroups of mapping class groups.

We will now develop a general theory of *stable injections* from right-angled Artin groups into mapping class groups, which is to say maps

$$A(\Gamma) \to \mathrm{Mod}(S)$$

which send vertices $v \mapsto \psi_v$ and which may not be injective, but which become injective after ψ_v is replaced by a sufficiently high power. In this subsection, we will prove the following result:

Theorem 5.6. *Let S be a surface and Γ a finite simplicial graph.*

1. *If Γ embeds as a subgraph of $\mathcal{C}(S)$ then $A(\Gamma)$ embeds as a subgroup of $\mathrm{Mod}(S)$.*
2. *If $A(\Gamma)$ embeds as a subgroup of $\mathrm{Mod}(S)$ then Γ embeds as a subgraph of $\mathcal{C}(S)_k$.*

It seems unlikely that one can make Theorem 5.6 an if and only if statement, in light of recent work of Casals–Ruiz, Duncan and Kazachkov (see [11]).

Corollary 5.7. *If Γ is a finite simplicial graph then there exists a surface $S = S(\Gamma)$ such that $A(\Gamma) < \mathrm{Mod}(S)$.*

Proof. It suffices to embed Γ in the curve complex of some surface. We leave the details as an exercise for the reader. □

In light of the Birman–Lubotzky–McCarthy bound on the size of abelian subgroups of mapping class groups, it is clear that there is no surface whose mapping class group contains every right-angled Artin group. Ranks of abelian subgroups do not provide the only obstruction to embedding right-angled Artin groups into mapping class groups. Recall that if Γ is triangle–free then any abelian subgroup of $A(\Gamma)$ has rank at most two:

Corollary 5.8. *For every surface S there exists a triangle–free graph $\Gamma = \Gamma(S)$ such that $A(\Gamma)$ does not embed in $\mathrm{Mod}(S)$.*

Proof. We have seen that $\mathcal{C}(S)$ (and hence $\mathcal{C}(S)_k$) have finite chromatic numbers. By a classical result of Erdös (see [17]), we have that there exist finite graphs which simultaneously have high girth and high chromatic number. Choosing a triangle–free graph with sufficiently high chromatic number, we see that no embedding of Γ in $\mathcal{C}(S)_k$ can exist. □

Theorem 5.6 has the following refinement, which underscores the stability of injections of right-angled Artin groups which we mentioned above:

Theorem 5.9. *Let*

$$\iota : \Gamma \to \mathcal{C}(S)$$

be an injection of graphs. Then for all sufficiently large N*, the map*

$$\iota_{*,N} : A(\Gamma) \to \mathrm{Mod}(S)$$

given by $v \mapsto T_{\iota(v)}^N$ *is injective.*

By $T_{\iota(v)}$ we mean the Dehn twist about the curve $\iota(v) \subset S$. We will prove a statement which is somewhat stronger than item (1) of Theorem 5.6 and even stronger than Theorem 5.9. Let $\{f_1, \ldots, f_k\}$ be mapping classes which are all either Dehn twists about simple closed curves or pseudo-Anosov on subsurfaces. We will say that these mapping classes form an *irredundant* collection if no pair of them generates a virtually cyclic subgroup of $\mathrm{Mod}(S)$. The *coincidence* graph of $\{f_1, \ldots, f_k\}$ has vertices $\{\mathcal{L}_1^+, \ldots, \mathcal{L}_k^+\}$, and edges whenever the corresponding laminations are disjoint. Observe that the coincidence graph is the same as the commutation graph in $\mathrm{Mod}(S)$. We will retain the terminology "coincidence graph" to underline the geometric nature of the result.

Theorem 5.10. *Let* $\{f_1, \ldots, f_k\}$ *be an irredundant collection of Dehn twists and pseudo-Anosov mapping classes on subsurfaces with coincidence graph* Γ*. There exists an* N *such that for all* $n \geq N$*, we have*

$$\langle f_1^n, \ldots, f_k^n \rangle \cong A(\Gamma).$$

We will see in the proof of Theorem 5.10 that there is a slight issue having to do with a pseudo-Anosov mapping class ψ on a subsurface S_0 which may or may not twist about the boundary ∂S_0, but we will not belabor this point here. The hypothesis one needs for the result to be true is that for each i, each lift of $\mathcal{L}_i^{\pm}$ must be dense in $\mathcal{L}_i^{\pm}$.

Part (1) of Theorem 5.6 follows from Theorem 5.10 in the case where $\{f_1, \ldots, f_k\}$ are all Dehn twists. We can now give a fairly quick proof of part (2) of Theorem 5.10, given part (1):

Proof of Theorem 5.6 part (2), assuming Theorem 5.10. Choose an embedding $\phi : A(\Gamma) \to \mathrm{Mod}(S)$. By replacing $\phi(v)$ by some positive power, we may assume that for each v the mapping class $\phi(v)$ is pure. Decompose $\phi(v)$ as a composition of commuting Dehn twists and pseudo-Anosov mapping classes, writing

$$\phi(v) = f_1^v \cdots f_{n(v)}^v.$$

There exists a collection of mapping classes C such that f_i^v is a power of some element in C for every vertex v of Γ and $1 \leq i \leq n(v)$. By choosing C with minimal cardinality, we may assume that no two elements of C generate a cyclic subgroup of $\mathrm{Mod}(S)$. Let X be the commutation graph of C. By Theorem 5.10, we have $\langle C \rangle < \mathrm{Mod}(S)$ is isomorphic to $A(X)$, again possibly after replacing each $\phi(v)$ by a further higher power and accordingly raising each element of C to a suitable power. We claim that X embeds in $\mathcal{C}(S)$, and then that Γ embeds in X_k, which will establish the result.

Write C as a union of powers of Dehn twists $\{g_1, \ldots, g_p\}$ about simple closed curves $\{\alpha_1, \ldots, \alpha_p\}$ and of pseudo-Anosov mapping classes $\{g_{p+1}, \ldots, g_{p+q}\}$ supported on subsurfaces $\{S_{p+1}, \ldots, S_{p+q}\}$. We let S_i be a sufficiently small regular neighborhood of α_i for $i \leq p$. The mapping classes g_i and g_j correspond to adjacent vertices of X if and only if S_i and S_j are disjoint. For each $i = p+1, p+2, \ldots, p+q$, we inductively choose an essential, nonperipheral simple closed curve β_i in the interior of S_i and apply a sufficiently large power of g_i to β_i to get α_i so that the following holds:

1. The curves α_i and α_j are disjoint if and only if S_i and α_j are as well, for $j = 1, 2, \ldots, i-1$;
2. The curve α_i and the surface S_j are disjoint if and only if the surfaces S_i and S_j are, for $j = i+1, i+2, \ldots, p+q$.

After this inductive process, we have that α_i and α_j are disjoint if and only if S_i and S_j are disjoint for $1 \leq i, j \leq p+q$. Furthermore, we can require the curves α_i and α_j to be noni-isotopic for $i \neq j$, even though S_i and S_j may be isotopic. Hence X coincides with the subgraph of $\mathcal{C}(S)$ induced by $Y = \{\alpha_1, \ldots, \alpha_{p+q}\}$.

Let ψ be an embedding from $A(\Gamma)$ into $A(X)$ such that $supp(\psi(v))$ is a clique for each vertex v of Γ and moreover,

$$\sum_{v \in V(\Gamma)} |supp(\psi(v))|$$

is minimal. Such a ψ exists, since the composition of $\phi : A(\Gamma) \to \langle C \rangle$ with the isomorphism $\langle C \rangle \cong A(X)$ is an embedding that sends each vertex of Γ to a word whose support is a clique. We claim that $supp(\psi(v)) \neq supp(\psi(v'))$ for distinct vertices v and v' of Γ. Suppose not, and write $\psi(v) = x_1^{p_1} \cdots x_k^{p_k}$ and $\psi(v') = x_1^{q_1} \cdots x_k^{q_k}$ such that $k > 0$,

$$\prod_i p_i \neq 0 \neq \prod_i q_i$$

and $\{x_1, \ldots, x_k\}$ span a clique in X. Note that for each vertex u of Γ adjacent to v, we have $[\psi(v), \psi(u)] = 1$ and so,

$$supp(\psi(v')) \cup supp(\psi(u)) = supp(\psi(v)) \cup supp(\psi(u))$$

spans a clique; this implies that $[\psi(v'), \psi(u)] = 1$ and v' is adjacent to u. Hence we have an automorphism $\xi : A(\Gamma) \to A(\Gamma)$ defined by $\xi(v) = vv'^{-1}$ and $\xi(x) = x$ for each $x \in V(\Gamma) \setminus \{v\}$. Let $\eta : A(\Gamma) \to A(\Gamma)$ be an embedding that maps v to v^{q_1} while fixing all the other vertices. Then

$$\psi' = \psi \circ \eta \circ \xi^{p_1} : A(\Gamma) \to A(X)$$

is an embedding such that

$$supp(\psi'(v)) \subseteq \{x_2, \ldots, x_k\} \subset supp(\psi(v))$$

and that $\psi'(y) = \psi(y)$ for $y \subseteq V(\Gamma) \setminus \{v\}$. This contradicts the minimality of supports assumption.

We now define a map $\delta : V(\Gamma) \to V(X_k)$ by $\delta(v) = supp(\psi(v))$. This map is an embedding by the preceding paragraph. Moreover, δ is easily seen to extend to a graph embedding $\Gamma \to X_k$ such that the $\delta(\Gamma)$ is an induced subgraph of X_k. □

We see that Theorem 5.10 has a number of hypotheses, all of which are necessary. Recall that inside of $F_2 \times F_2$ we produced a finitely generated group G which is not finitely presented. We can embed $F_2 \times F_2$ inside of the mapping class group of a surface of genus two or more by taking four simple closed curves $\{a, b, c, d\}$ such that a and c intersect and such that b and d intersect, and all other pairs are disjoint. The Dehn twists about these curves may not generate a copy of $F_2 \times F_2$, but their squares will. Thus, we can embed the group G as a subgroup of the mapping class group in such a way that G is generated by two squares of Dehn twists about disjoint simple closed curves and the square of a twist about a multicurve with two components. Replacing these generators by any nonzero power results in a subgroup of the mapping class group which is still isomorphic to G. An identical example can be produced using pseudo-Anosov mapping classes on subsurfaces in lieu of Dehn twists. Thus, the assumption that $\{f_1, \ldots, f_k\}$ be Dehn twists or pseudo-Anosov on subsurfaces is essential.

One might venture to guess that one N should work for all collections $\{f_1, \ldots, f_k\}$, so that the power of the mapping classes one must take would depend only on the underlying surface and not on the mapping classes. This is not true for a somewhat silly reason. Let $\{c_1, c_2\}$ be intersecting curves and let T_1, T_2 be the corresponding Dehn twists. Let N be a purported minimal power which should make Theorem 5.10 true for all finite collections of mapping classes. Let $c_3 = T_1^N(c_2)$ and let T_3 be the Dehn twist about c_3. Then according to the result, we have

$$\langle T_1^N, T_2^N, T_3^N \rangle \cong F_3.$$

However,

$$T_3^N = T_{T_1^N(c_2)} = T_1^N T_2^N T_1^{-N},$$

a contradiction since T_1^N, T_2^N and T_3^N are all N^{th} powers of Dehn twists about simple closed curves.

In order to prove Theorem 5.10, we will collect some facts from hyperbolic geometry which we will need. For the rest of the proof, we will be fixing a finite volume complete hyperbolic metric on S, and all laminations will be assumed to be geodesic and compactly supported. In order to apply mapping classes to geodesics, we will choose homeomorphism representatives of mapping classes, apply them to geodesics, and straighten out the resulting curves to be geodesics.

If γ_1 and γ_2 are distinct geodesics on S that intersect at a point $x \in S$, we will measure the *angle of intersection* $\alpha_x(\gamma_1, \gamma_2)$ between γ_1 and γ_2 at x as the smaller of the two angles cut out by γ_1 and γ_2. Thus, $\alpha_x(\gamma_1, \gamma_2) \in (0, \pi/2]$. We will suppress the subscript if there is no confusion possible concerning the point of intersection. We will say that two geodesics intersects with a given angle if there exists an intersection point between those two geodesics with that angle of intersection.

Lemma 5.11. *Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be distinct geodesic laminations such that each leaf of $\mathcal{L}_i$ is dense in $\mathcal{L}_i$. Then there exists an $\epsilon > 0$ which depends only on $\mathcal{L}_1$ and $\mathcal{L}_2$ such that any two leaves $\gamma_1 \subset \mathcal{L}_1$ and $\gamma_2 \subset \mathcal{L}_2$ satisfy $\alpha(\gamma_1, \gamma_2) > \epsilon$ at any intersection point of γ_1 and γ_2.*

Proof. Suppose the contrary so that there exists a sequence of points $\{x_i\} \subset \mathcal{L}_1$ such that

$$\alpha_{x_i}(\mathcal{L}_1, \mathcal{L}_2) \to 0.$$

Here, the intersection pairing denotes the intersection angle of the unique leaves of $\mathcal{L}_1$ and $\mathcal{L}_2$ passing through x_i. By the compactness of $\mathcal{L}_1$ and $\mathcal{L}_2$, the points $\{x_i\}$ accumulate at a point $x \in \mathcal{L}_1 \cap \mathcal{L}_2$. It follows that at x, the laminations $\mathcal{L}_1$ and $\mathcal{L}_2$ share a leaf, since a geodesic on S is determined up to equality by its direction and a point through which it passes. This leaf is dense in both $\mathcal{L}_1$ and in $\mathcal{L}_2$ so that $\mathcal{L}_1 = \mathcal{L}_2$, a contradiction. □

For the next three lemmas, we will assume that $\mathcal{L}_1$ and $\mathcal{L}_2$ are limiting laminations of mapping classes f_1 and f_2 which are Dehn twists or pseudo-Anosov mapping classes on subsurfaces, with $\mathcal{L}_1 \neq \mathcal{L}_2$.

Lemma 5.12. *Suppose $\mathcal{L}_1 \cap \mathcal{L}_2 \neq \emptyset$, and write $\epsilon > 0$ for the minimal angle of intersection between two leaves of $\mathcal{L}_1$ and $\mathcal{L}_2$. Then there is a $\delta > 0$ such that if γ is any geodesic which intersects a leaf of $\mathcal{L}_1$ with angle at most δ, then γ intersects $\mathcal{L}_2$ with angle at least $\epsilon/2$.*

Proof. If $\mathcal{L}_1$ consists of a single closed leaf then the conclusion is obvious, since if δ is small enough then γ will diverge from $\mathcal{L}_1$ very slowly at first.

When $\mathcal{L}_1$ has no closed leaves, consider an interval I (with nonempty interior) which is transverse to $\mathcal{L}_1$. A standard fact about pseudo-Anosov laminations is that for any intersection point of I and $\mathcal{L}_1$, following the leaf of $\mathcal{L}_1$ through that intersection point will return to a prescribed subinterval $I_0 \subset I$ in finite time. We set I_0 to be the subinterval of I for which leaves passing through I_0 intersect $\mathcal{L}_2$ before returning to I. Thus, if γ intersects a leaf of $\mathcal{L}_1$ within I with a sufficiently small angle δ then γ will diverge very slowly from that leaf. The compactness of $\mathcal{L}_1$ allows us to choose a positive δ uniformly over $\mathcal{L}_1$. □

Lemma 5.13. *Let $\delta > 0$ be fixed and let γ be a geodesic which intersects a leaf of $\mathcal{L}_1^-$ with angle at least $\epsilon > 0$. Then there is an N which depends only on ϵ, δ and f_1 such that for all $n \geq N$, we have that $\alpha(f_1^n(\gamma), \mathcal{L}_1^+) < \delta$.*

Proof. Exercise. □

We reiterate that in the previous lemma, we mean that there exists an intersection point between $f_1^n(\gamma)$ and $\mathcal{L}_1^+$ where the angle is less than δ. The important content of the previous lemma is that the N is independent of γ and of the other potential intersection points of γ and $\mathcal{L}_1^+$. Intuitively, the lemma says that if γ is not nearly parallel to a leaf of $\mathcal{L}_1$ at an intersection point, then for any sufficiently high power of f_1, the resulting curve $f_1^n(\gamma)$ will be nearly parallel to a leaf of $\mathcal{L}_1^+$ in a way which is independent of the choice of γ. Of course the distinction between $\mathcal{L}_1^+$ and $\mathcal{L}_1^-$ is irrelevant in the case of a Dehn twist, but the distinction is important when f_1 is pseudo-Anosov on a subsurface.

Lemma 5.14. *Suppose $\mathcal{L}_1$ and $\mathcal{L}_2$ are disjoint and let $\epsilon > 0$ be fixed. Let γ be a geodesic which intersects both $\mathcal{L}_1$ and $\mathcal{L}_2$. There exists an $\delta > 0$ which depends only on ϵ, f_1 and f_2 with the following property: if there exists an intersection point between $\mathcal{L}_1$ and γ with angle at most δ then for all N, we have $\alpha(f_2^N(\gamma), \mathcal{L}_1) < \epsilon$.*

Proof. Exercise. □

Again, we mean that there exists an intersection point between $f_2^N(\gamma)$ and $\mathcal{L}_1$ with angle at most ϵ. Intuitively, the lemma is saying that if γ is nearly parallel to a leaf of $\mathcal{L}_1$ then no matter what power of f_2 we apply to γ, the resulting geodesic will still be nearly parallel to a leaf of $\mathcal{L}_1$.

The proof of Theorem 5.10 proceeds by induction on a certain type of length of words in $A(\Gamma)$ which is related, but not quite equal to the word length. Let $1 \neq w \in A(\Gamma)$ be a reduced word. Recall that we can write w as a product of *central words*, which is to say as a product of subwords $w_k \cdots w_1$ where the support of each w_i is contained in a clique of Γ. We will write w in *left–greedy form* as follows. Write w as a product of central words with k minimal. If $k = 1$ then we declare w to be in left–greedy form. Suppose $k \geq 2$. If $v \in supp(w_{k-1})$ commutes with w_k then we move each occurrence of v to w_k. We repeat this left–shifting with for each vertex in the support of w_{k-1}. We then repeat this process for vertices in the support of w_{k-2}, shifting vertices which commute with w_{k-1} to w_{k-1}, and then further shifting vertices which commute with w_k to w_k. Repeating this process for all w_i with $i < k$ will terminate in finite time, since we are always shifting letters to the left and since there are at most k central words in the expansion of w at any stage. The result is a reduced word $w = w_k \cdots w_1$ with the property that each w_i is a central word and that for each $i < k$ and each vertex $v \in supp(w_i)$, there exists a vertex in $supp(w_{i+1})$ which does not commute with v.

The proof of Theorem 5.10 we give now is more or less complete. The precise estimates can be worked out by the reader, or the reader may consult [32].

Proof of Theorem 5.10. For brevity, we will assume that the mapping classes under consideration are all Dehn twists. The case where pseudo-Anosov mapping classes on subsurfaces are allowed carries over without difficulty. We are given Dehn twists $\{f_1, \ldots, f_k\}$ with limiting laminations $\{\mathcal{L}_1, \ldots, \mathcal{L}_k\}$. There is a positive constant c such that when $i \neq j$, the minimal angle of intersection between $\mathcal{L}_i$ and $\mathcal{L}_j$ is at least c. Furthermore, there is a δ such that if $\mathcal{L}_i \cap \mathcal{L}_j \neq \emptyset$ and if γ is any geodesic which intersects $\mathcal{L}_i$ at an angle of at most δ then γ intersects $\mathcal{L}_j$ at an angle of at least $c/2$.

For any ϵ, we have that there is an M such that if γ intersects $\mathcal{L}_i$ at an angle of at least $c/2$ then $f_i^n(\gamma)$ intersects $\mathcal{L}_i$ at an angle less than ϵ for all $n \geq M$. We choose an ϵ (and corresponding M) so that if γ intersects $\mathcal{L}_i$ at an angle of at most ϵ and if F is a composition of Dehn twists on a multicurve consisting of components of $\{\mathcal{L}_1, \ldots, \mathcal{L}_k\}$ which are disjoint from $\mathcal{L}_i$ then

$$\alpha(F(\gamma), \mathcal{L}_i) < \delta.$$

We will let γ be a simple closed geodesic in S which intersects each $\mathcal{L}_i$. By choosing ϵ and δ smaller if necessary, we may assume that γ intersects each $\mathcal{L}_i$ with angle exceeding both ϵ and δ at each intersection point. We let M' be a sufficiently large integer, so that $\alpha(f_i^n(\gamma), \mathcal{L}_i) < \epsilon$ for each $n \geq M'$. We set $N = \max\{M, M'\}$, which is the N we wish to produce in the statement of the theorem.

Let $1 \neq w \in A(\Gamma)$ be a word, written in left–greedy form as $w_k \cdots w_1$. We choose vertices $v_1, \ldots, v_k$ such that v_i is in the support of w_i and such that v_i and v_{i+1} do not commute for $i < k$. By abuse of notation, we identify v_i with f_i^N. We claim that for each i, the curve

$$w_i \cdots w_1(\gamma)$$

has an intersection point with $\mathcal{L}_i$ with angle at most δ. Since each angle of intersection between γ and $\mathcal{L}_i$ greater than δ, it follows that

$$w_i \cdots w_1(\gamma) \neq \gamma,$$

whence $w_i \cdots w_1$ represents a nontrivial mapping class. Thus, the claim will establish the result.

Let us apply w_1 to γ, applying all the copies of v_1 first. Applying any nonzero power of v_1 to γ results in a curve which intersects $\mathcal{L}_1$ with angle at most ϵ. Applying the rest of the letters of w_1 may increase the angle of intersection with $\mathcal{L}_1$, but not past δ.

Now consider

$$\gamma_i = w_i \cdots w_1(\gamma).$$

We have that γ_i intersects $\mathcal{L}_{i+1}$ in at least one point with angle $c/2$ since there is a point of intersection between γ_i and $\mathcal{L}_i$ with angle at most δ. Apply w_{i+1} to γ_i, applying all the copies of v_{i+1} first. The resulting curve has an angle of intersection at most ϵ with $\mathcal{L}_{i+1}$. Applying the rest of w_{i+1} to the resulting curve may increase the minimal angle of intersection, but not past δ. The claim follows by induction on k. □

5.7. **Mapping class group subgroups of right-angled Artin groups.** A natural question which arises in light of the results of the previous sections is whether one can reverse the directions of the arrows and find embeddings of mapping class groups inside of right-angled Artin groups. Since mapping class groups usually have torsion and right-angled Artin groups are torsion–free, mapping class groups themselves cannot be embedded in right-angled Artin groups. However, mapping class groups are virtually torsion–free, so one can wonder whether mapping class groups virtually embed in right-angled Artin groups. We will show that generally they do not:

Theorem 5.15. *Let $S = S_{g,n}$ where $g \geq 3$ or $g = 2$ and $n \geq 2$. Then no finite index subgroup of* $\mathrm{Mod}(S)$ *occurs as a subgroup of a right-angled Artin group.*

Results of this ilk have been proved by other authors, such as Kapovich and Leeb in [27]. The idea behind the proof of Theorem 5.15 is that under the assumptions on S the group $\mathrm{Mod}(S)$ contains a subgroup which is the fundamental group of a nontrivial circle bundle over a closed surface with trivial monodromy. Then, one shows that right-angled Artin groups have a certain residual property which is inherited by subgroups and which does not hold for nontrivial circle bundles over closed surfaces.

Lemma 5.16. *Let S be as in the statement of Theorem 5.15 and let $G = \pi_1(T^1 S_2)$, the fundamental group of the unit tangent bundle of a genus two closed surface. Then* $\mathrm{Mod}(S)$ *contains a copy of G.*

Proof. This follows from a version of the Birman exact sequence (see [21]). □

A group G is called *residually finite rationally solvable*, abbreviated RFRS, if there is a sequence of finite index, nested, normal subgroups $\{G_i\}$ of G satisfying:

(1) Residual finiteness:

$$\bigcap_i G_i = \{1\}.$$

(2) Rational solvability:

$$\ker\{G_i \to H_1(G_i, \mathbb{Q})\} < G_{i+1}.$$

Lemma 5.17. *Let G be RFRS and let $H < G$ be a subgroup. Then H is RFRS.*

Proof. Let $\{G_i\}$ be a tower of subgroups witnessing the fact that G is RFRS, and let $H_i = H \cap G_i$. Evidently, we have

$$\bigcap_i H_i = \{1\}.$$

Clearly we have H_i/H_{i+1} is abelian, so that $\ker\{H_i \to H_1(H_i, \mathbb{Z})\} < H_{i+1}$. Let $h \in H_i$ be an element whose image in $H_1(H_i, \mathbb{Z})$ is nontrivial and torsion. Then there is an N such that h^N is a product of commutators in H_i, so that h^N is a product of commutators in G_i. Since $h \in G_i \setminus G_{i+1}$, we have that h itself is not a product of commutators in G_i (since G_i/G_{i+1} is abelian), whence h is forced to lie in G_{i+1}, a contradiction. □

The following lemma is due to Agol in [1], where the RFRS condition was originally defined. We more or less follow Agol's proof. Observe that RFRS groups must be torsion–free, so a group with torsion can only hope to be virtually RFRS.

Lemma 5.18. *Right-angled Coxeter groups are virtually RFRS. In particular, right-angled Artin groups are virtually RFRS.*

Proof. Consider the right-angled Coxeter group $C(\Gamma)$, where Γ has n vertices. Using the Davis–Januszkiewicz construction, we can define a mirror structure on a union of coordinate cubes inside of the unit cube in $\mathbb{R}^n$ in order to obtain a complex U on which $C(\Gamma)$ acts with compact quotient Q. The quotient could be reasonably called an orbifold, though it is only an orbifold in the cube complex sense and not necessarily in the manifold sense.

The action of $C(\Gamma)$ on U is by reflection across certain distinguished hyperplanes, all of which meet at right angles, and which project to a finite collection $\{H_1, \ldots, H_k\}$ of dual hypersurfaces of Q. Taking the cover $\widetilde{Q}$ corresponding to the abelianization $(\mathbb{Z}/2\mathbb{Z})^n$ of $C(\Gamma)$, the hypersurfaces $\{H_1, \ldots, H_k\}$ lift to orientable, two–sided, embedded dual hypersurfaces of $\widetilde{Q}$.

We take a cofinal sequence of 2–fold covers

$$\cdots \to Q_2 \to Q_1 \to Q_0 = Q$$

which are given by reflection in the hypersurfaces $\{H_1, \ldots, H_k\}$. Each Q_i induces a cover of $\widetilde{Q}_i$ by taking the cover corresponding to

$$G_i = \pi_1(Q_i) \cap [C(\Gamma), C(\Gamma)] = \pi_1(Q_i) \cap \pi_1(\widetilde{Q}).$$

We claim that $\{G_i\}$ witnesses the fact that $C(\Gamma)$ is virtually RFRS. We have that

$$\bigcap_i G_i = \{1\}$$

since we are taking a cofinal sequence of covers of Q. Furthermore, note that G_i/G_{i+1} is either trivial or $\mathbb{Z}/2\mathbb{Z}$. We just need to show that the map $G_i \to G_i/G_{i+1}$ factors through $H_1(G_i, \mathbb{Z})/torsion$.

Let $g \in G_i \setminus G_{i+1}$. We represent g by a closed loop γ in $\widetilde{Q}_i$ and project it to Q_i. Generically, γ travels around Q_i, bouncing off of the various dual hypersurfaces. If γ bounces off each hypersurface an even number of times then γ lifts to Q_{i+1}, so we may assume there is at least one dual hypersurface which γ bounces off an odd number of times. The intersection pairing with that hypersurface induces a surjective homomorphism from $\pi_1(\widetilde{Q}_i) \to \mathbb{Z}$, since the hypersurface is embedded and two–sided and nonseparating (since γ is closed and intersects it an odd number of times). It follows that g is sent to an infinite order element in $H_1(G_i, \mathbb{Z})$, which is what we needed to show. □

By the Davis–Januszkiewicz construction, each right-angled Artin group is commensurable with a right-angled Coxeter group, whence right-angled Artin groups are virtually RFRS.

Lemma 5.19. *Let*

$$S^1 \to M \to S$$

be a circle bundle over a closed, orientable surface S with trivial monodromy and with nontrivial Euler class. Then $\pi_1(M)$ is not virtually RFRS.

Proof. Let $\widetilde{M} \to M$ be a finite cover. Then $\widetilde{M}$ is still a circle bundle over a closed orientable surface with trivial monodromy. By general cohomology of groups, we have that the Euler class of the bundle $\widetilde{M}$ is nonzero (this is a consequence of the fact that if G is a group and H is a subgroup of index n then the composition of the restriction with the corestriction homomorphism is just multiplication by n on $H^*(G, \mathbb{Z})$. The reader may consult [10]). Let $G = \pi_1(\widetilde{M})$ and write $\{G_i\}$ for the tower of subgroups witnessing the fact that G is RFRS. We have a short exact sequence

$$1 \to \mathbb{Z} \to G \to \overline{G} \to 1,$$

where $\overline{G}$ is the fundamental group of the base space of the fibration and the leftmost copy of $\mathbb{Z}$ is central in G. The fact that the Euler class of this central extension is nonzero implies that some nontrivial element of the central copy of $\mathbb{Z}$ is a product of commutators in G. It follows that this copy of $\mathbb{Z}$ has torsion image in $H_1(G, \mathbb{Z})$, so that it is contained in G_2.

Since G_i has finite index in G, the group G_i also has a description as a central extension of a closed surface group by $\mathbb{Z}$ with a nontrivial Euler class. It follows that this central copy of $\mathbb{Z}$ satisfies

$$\mathbb{Z} < \bigcap_i G_i,$$

a contradiction. □

Since the unit tangent bundle of a genus two surface has a nontrivial Euler class, Theorem 5.15 follows immediately.

6. Right-angled Artin groups in mapping class groups and the isomorphism problem for finitely presented subgroups

In this section, we would like to discuss a consequence of the fact that each right-angled Artin group occurs in some mapping class group. We will discuss Bridson's recent result on decision theory in mapping class groups, namely:

Theorem 6.1 (See [9]). *Let $\mathcal{FP}(S)$ be the class of finitely presented subgroups of* $\mathrm{Mod}(S)$*, where $\chi(S) \ll 0$. There exists no algorithm which decides whether two members of $\mathcal{FP}(S)$ are isomorphic.*

This result is really a fact about right-angled Artin groups, combined with effective embeddings of right-angled Artin groups into mapping class groups given by Theorem 5.10, for example (also [13], also [14], [15], [12]). We will not give a self–contained proof of Theorem 6.1, but we will give references for proofs of relevant results.

6.1. **A criterion for unsolvability of the isomorphism problem.** For this subseciton, we will follow [8] closely. In [8], Bridson and Miller give a method of constructing families of finitely presented subgroups of an ambient finitely presented group which have an unsolvable isomorphism problem. Specifically, they prove the following result:

Theorem 6.2. *Let F be an nonabelian free group and let*

$$1 \to K \to G \to L \to 1$$

be a short exact sequence of groups and suppose:

1. *G is torsion–free and hyperbolic.*
2. *K is infinite and finitely generated.*
3. *L is a nonabelian free group.*

Then there is a recursive collection $\{\Delta_i\}_{i\in\mathbb{N}}$ *of subsets of* $G \times G \times F$ *together with finite presentations* $\{\langle \Delta_i \mid R_i \rangle\}$ *of the subgroups they generate such that no algorithm can determine whether*

$$\langle \Delta_i \mid R_i \rangle \cong \langle \Delta_0 \mid R_0 \rangle.$$

Recall that a group is called *hyperbolic* if its Cayley graph is a δ–hyperbolic metric space for some $\delta \geq 0$. Recall that this means that if $[x, y] \cup [y, z] \cup [x, z]$ is a geodesic triangle then $[x, z]$ is contained in a δ–neighborhood of $[x, y] \cup [y, z]$. An exposition on some of the properties of hyperbolic groups can be found in [23], for example.

Note that the finite presentability of the subgroups generated by the subsets $\{\Delta_i\}$ is implicit in the theorem. In this subsection, we will discuss the ideas that go into Theorem 6.2 in order to decrease the number of "black boxes" used in the proof of Theorem 6.1.

Let $A \times B$ be a direct product of two groups. A subgroup $H < A \times B$ is called a *subdirect product* if the projection map from $A \times B$ to each direct factor is a surjection when restricted to H. The following is given in [8] as a criterion for the finite presentability of certain semidirect products of products of finitely generated groups by subdirect products of free groups:

Proposition 6.3. *Let* K_1 *and* K_2 *be finitely generated groups and let* L_1 *and* L_2 *be finitely generated free groups. Let* $\phi_i : L_i \to \mathrm{Aut}(K_i)$ *be homomorphisms such that the associated semidirect products of* K_i *with* L_i *are finitely presented.*

Furthermore, let $p : F \to H$ *be surjective a homomorphism from a finitely presented group* F *to a subdirect product* $H < L_1 \times L_2$*. Then the associated semidirect product of* $K_1 \times K_2$ *by* F*, with* F *acting by* $(\phi_1, \phi_2) \circ p$*, is finitely presented.*

Observe that we do not assume K_1 and K_2 to be finitely presented. In [8], the surjectivity of p is not mentioned in the hypotheses of the proposition though it appears to be essential and is implicit in the proof: for instance, let K_1 be a non–finitely presented group and let K_2 and F both be the trivial group. Then $K_1 \times K_2 \cong K_1$ and the semidirect product of this group with F is isomorphic to K_1.

Proof of Proposition 6.3. For $i = 1, 2$, let $\langle B_i \mid T_i \rangle$ be a presentation for K_i with B_i finite and let C_i be a free basis for L_i. Let

$$R_i = \{c^{-1}bc\beta^{-1}\},$$

where $c \in C_i$, $b \in B_i$ and $\beta = \phi_i(c)(b)$. The presentation $\langle B_i, C_i \mid T_i, R_i \rangle$ is a presentation of the semidirect product of K_i with L_i. Since this semidirect product is finitely presented by assumption, there is a presentation of the form $\langle B_i, C_i \mid S_i, R_i \rangle$, where $S_i \subset T_i$ is finite. To see this, let T_i be the set of words in B_i which represent the identity in K_i and w be a word in B_i and C_i. We may use relations in R_i to rewrite w as an equivalent word $w_1 \cdot w_2$, where w_1 is a word in C_i and w_2 is a word in B_i. If w_1 is reduced and nontrivial then w does not represent the identity. If w_1 is trivial then w represents the identity if and only if w_2 is trivial in K_i. Thus, every word which represents the identity in the semidirect product can be written as elements of T_i, since T_i consisted of all words in K_i which represent the identity. Thus, such an S_i exists.

Writing Z for the set of commutators $[b_1, b_2]$ with $b_i \in B_i$ makes $\langle B_1, B_2 \mid T_1, T_2, Z \rangle$ a presentation for $K_1 \times K_2$. Let $F = \langle X \mid Y \rangle$. We will write

$$U_i = xbx\gamma^{-1},$$

where $b \in B_i$, $x \in X$, and where $\gamma = \phi_i \circ p(x)(b)$. We have that

$$\langle B_1, B_2, X \mid T_1, T_2, Z, Y, U_1, U_2 \rangle$$

presents the semidirect product of $K_1 \times K_2$ by F.

For each $c \in C_i$ there is a word ρ in X which maps to c under the composition

$$F \to L_1 \times L_2 \to L_i$$

by the assumption of the surjectivity of p. We write V_i for the set of words $c\rho^{-1}$, which is a finite set. Note that then

$$\langle B_1, B_2, C_1, C_2, X \mid T_1, T_2, Z, U_1, U_2, V_1, V_2\rangle$$

presents the semidirect product of $K_1 \times K_2$ by F. We have that the sets R_i for $i = 1, 2$ are consequences of the relations in the previous presentation. We can add them to the presentation to get

$$\langle B_1, B_2, C_1, C_2, X \mid T_1, T_2, R_1, R_2, Z, U_1, U_2, V_1, V_2\rangle.$$

Recall that T_i is a consequence of finitely many relations $S_i \cup R_i$, so that

$$\langle B_1, B_2, C_1, C_2, X \mid S_1, S_2, R_1, R_2, Z, U_1, U_2, V_1, V_2\rangle$$

presents the semidirect product of $K_1 \times K_2$ by F. All the sets on the right are finite, so the presentation is finite. □

A general method for showing that a particular decision problem in group theory is undecidable is by showing it is equivalent to another decision problem which is known to be undecidable. We will be using the undecidability of the generation problem for products of nonabelian free groups as shown by C. Miller in [37]:

Theorem 6.4. *Let L be a nonabelian free group and X a finite generating set for $L \times L$. For any sufficiently large M, there exists a recursive sequence Λ_n of finite sets of words in X such that*

1. *Each Λ_n has cardinality M.*
2. *Each of the subgroups $\Lambda_n \subset L \times L$ is a subdirect product.*
3. *$\langle\Lambda_0\rangle = L \times L$.*
4. *There is no algorithm which decides whether or not $\langle\Lambda_n\rangle = L \times L$.*
5. *If $\langle\Lambda_n\rangle \neq L \times L$ then $\langle\Lambda_n\rangle$ is not finitely presented.*

Let us discuss the sets Λ_n of Theorem 6.4 in some more detail. Let $\{c_1, \ldots, c_k\}$ be a free basis for L. One constructs a sequence of quotients of L of the form

$$Q_n = \langle c_1, \ldots, c_k \mid q_{n,1}, \ldots, q_{n,m-k}\rangle$$

which have the property that deciding whether Q_n is infinite or trivial is unsolvable. Then we set

$$\Lambda_n = \{(c_1, c_1), \ldots, (c_k, c_k), (q_{n,1}, 1), \ldots, (q_{n,m-k}, 1)\}.$$

Evidently, Λ_n generates all of $L \times L$ if and only if Q_n is trivial.

Before we outline the proof of Theorem 6.2, we need to understand the uniqueness of certain direct product decompositions. If a group decomposes as a direct product of two groups, the decomposition may not be unique: consider $\mathbb{Z}^2$ for instance. For products of non–elementary (not virtually cyclic), torsion–free hyperbolic groups however, direct product decompositions are unique.

Let G be a torsion–free hyperbolic group. The centralizer of each nontrivial element g of G is cyclic and is generated by the unique element h such that $h^k = g$, where $k > 0$ is maximal. Write $D = G_1 \times \cdots \times G_n$ for a direct product of non–elementary torsion–free hyperbolic groups. Let $x, y \in D$ be non–commuting elements, written as $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_n)$. So, we have x_i and y_i do not commute for some i. If $z = (z_1, \ldots, z_n)$ commutes with both x and y, we have z_i is trivial.

A subgroup $H < G_1 \times \cdots \times G_n$ is *full* if $H \cap G_i$ contains a non–commuting pair x_i, y_i for each i. Observe that the set of centralizers of non–commuting elements in H has a finite set of maximal elements, namely

$$M_i = H \cap (G_1 \times \cdots \times G_{i-1} \times \{1\} \times G_{i+1} \times \cdots \times G_n)$$

for each i. It is easy to check that $\{M_1, \ldots, M_n\}$ are sent to maximal centralizers of non–commuting pairs of elements under any isomorphism of groups. It follows that $\{M_1, \ldots, M_k\}$ is a characteristic collection of subgroups of H. By looking at intersections of $n - 1$ distinct subgroups among $\{M_1, \ldots, M_k\}$, we see that $H \cap G_i$ is characteristic in $G_1 \times \cdots \times G_n$. Thus:

Proposition 6.5. *Let H be a full subgroup of $D = G_1 \times \cdots \times G_n$. Then $\{H \cap G_i\}$ is a characteristic collection of subgroups of H, as is*

$$(H \cap G_1) \times \cdots \times (H \cap G_n).$$

It follows that the decomposition of D as a direct product of groups is unique.

Proof of Theorem 6.2. Let M and $\{\Lambda_n\}$ be as in Theorem 6.4. Let F be a free group of rank $2M$ with basis

$$\{x_1, \ldots, x_M, y_1, \ldots, y_M\}.$$

Write $\Lambda_n = \{\lambda_{n,1}, \ldots, \lambda_{n,m}\}$. Let $p_n : F \to L \times L$ be given by $p_n(x_i) = \lambda_{n,i}$ and $p_n(y_i) = 1$.

Let $K = K_i$ and $L = L_i$ as in Proposition 6.3. Let G be the semidirect product of K_i by L_i and let E_n be the semidirect product of $K \times K$ by F induced by p_n and the action of L_i on K_i. Suppose that K is generated by $\{b_1, \ldots, b_p\}$. Note that the group E_n can be viewed as the subgroup of $G \times G \times F$ generated by $\{(b_i, 1, 1)\}$ for all i, $\{(1, b_i, 1)\}$ for all i, $\{(c_i, c_i, x_i)\}$ for $i \leq k$, $\{(q_{n,j}, 1, x_{k+j})\}$ for $j \leq m - k$, and by $\{(1, 1, y_i)\}$ for all i. Call this collection of elements Δ_n.

Proposition 6.3 implies that E_n is finitely presented since $\langle \Lambda_n \rangle$ is a subdirect product, so that we have $E_n = \langle \Delta_n \mid R_n \rangle$ as in the statement of the result.

We claim that $E_n \cong E_0$ if and only if Λ_n generates all of $L \times L$, which will show that the isomorphism problem for the collection $\{E_n\}$ is unsolvable.

Suppose $\langle \Lambda_n \rangle \cong L \times L$. It suffices to find an automorphism ψ of F_{2M} such that $p_0 = p_n \circ \psi$. This is guaranteed by Bridson's refinement of a result of Rapaport ([38]), which says that if H is any group and if $\phi_1, \phi_2 : F_{2M} \to H$ are maps with equal images and such that $\phi_i(y_j) = 1$ for $i = 1, 2$ and $i \leq M$ then there exists a $\psi \in \mathrm{Aut}(F_{2M})$ which intertwines ϕ_1 and ϕ_2.

For the converse, note that K is infinite and normal in G which is torsion–free and hyperbolic. The fact that G surjects to a nonabelian free group implies that G is non–elementary. It follows that K is nonabelian and hence contains a pair of non–commuting elements. In $G \times G \times F$, the intersection of $\langle \Delta_n \rangle$ with each of the first two factors is just a copy of K and is hence a finitely generated group. Label these intersections by K_1 and K_2 respectively.

The intersection of $\langle \Delta_n \rangle$ with the third factor is just the kernel of the map $p_n : F \to L \times L$. Observe that the assumptions on Λ_n imply that $\langle \Lambda_n \rangle$ is finitely presented if and only if $\langle \Lambda_n \rangle = L \times L$. Thus the kernel of p_n is finitely normally generated if and only if $p_n(F) = L \times L$.

Proposition 6.5 implies that $\{K_1, K_2, F \cap \langle \Delta_n \rangle\}$ is a characteristic collection of subgroups of $\langle \Delta_n \rangle$, which are finitely normally generated if and only if $\langle \Delta_n \rangle = L \times L$. Thus if Λ_n does not generate all of $L \times L$ then $\langle \Delta_n \mid R_n \rangle$ is not isomorphic to $\langle \Delta_0 \mid R_0 \rangle$.

This completes the proof. Indeed, suppose there exists an algorithm which determines whether or not $\langle \Delta_n \mid R_n \rangle \cong \langle \Delta_0 \mid R_0 \rangle$. Then the values of n for which the algorithm returns "yes" certainly includes the n for which $p_n(F) = L \times L$. It cannot include any values for which $p_n(F) \neq L \times L$ since we have argued that then $\langle \Delta_n \mid R_n \rangle$ is not isomorphic to $\langle \Delta_0 \mid R_0 \rangle$. Thus, the existence of such an algorithm contradicts Theorem 6.4. □

6.2. **The Rips construction.** If one could effectively embed such a group G as in Theorem 6.2 in a right-angled Artin group then one could immediately conclude that the isomorphism problem for finitely presented subgroups of certain right-angled Artin groups (and by Theorem 5.10 for all sufficiently complicated mapping class groups) is unsolvable.

The essential idea in constructing such a G comes from a now classical construction due to Rips (see [39]). We first recall some notions from small cancellation theory. Let $G = \langle F \mid R \rangle$ be a presentation for a group G. We will assume that R is *symmetrized*, which is to say that R is closed under taking cyclic permutations of words in R and under taking inverses. We will call a reduced word u a *piece* for this presentation of G if it is a maximal initial segment of two distinct elements of R.

Let $0 < \lambda < 1$ and let $G = \langle F \mid R \rangle$ be a symmetrized presentation for G. We say that G satisfies $C'(\lambda)$, called a *metric small cancellation condition*, if whenever u is a piece for the given presentation of G and u is a subword of some $r \in R$ then $|u| \leq \lambda |r|$. A classical fact about small cancellation is the following result:

Proposition 6.6. *Let G be a finitely presented group which satisfies the $C'(1/6)$ metric small cancellation condition. Then G is hyperbolic. If all the elements of R are not proper powers then G is torsion–free.*

The proof of the previous proposition follows more or less from the following result, which is often called Greedlinger's Lemma:

Proposition 6.7. *Let G satisfy the $C'(\lambda)$ metric small cancellation condition for some $\lambda \leq 1/6$, and let w be a reduced word which represents the identity in G. Then there exists a subword v of w and a relator of G such that v is also a subword of r and such that*

$$|v| > (1 - 3\lambda)|r|.$$

Greedlinger's Lemma allows us to solve the word problem for $C'(\lambda)$ groups using Dehn's algorithm. Precisely, we take a reduced word w and check its subwords of w against subwords of the relators r of G which have length more than $1/2|r|$. It we find such a subword v, we can write $r = v \cdot u$ (since the presentation for G is symmetrized) and replace v by u^{-1} in w. After reducing if necessary, we obtain another word which has strictly shorted length than w and which represents the same element of G. Greedlinger's Lemma implies that if w represents the identity in G then such a subword v always exists. When Dehn's algorithm can be used to solve the word problem in a group G, one can show that the group satisfies a linear isoperimetric inequality and is hence hyperbolic (see [23]).

The following result is often called the *Rips construction*:

Theorem 6.8. *Let $0 < \lambda < 1$ and let Q be a finitely presented group. There exists a short exact sequence of groups*

$$1 \to K \to G \to Q \to 1$$

such that

1. *G is finitely presented and satisfies the $C'(\lambda)$ metric small cancellation condition.*
2. *K is finitely generated.*

Proof. We will follow Rips' original proof. Let Q be presented as

$$Q = \langle a_1, \ldots, a_m \mid r_1, \ldots, r_n \rangle.$$

We let G be generated by $\{a_1, \ldots, a_m, b_1, b_2\}$, and we impose relations

$$r_i b_1 b_2^{r_i} b_1 b_2^{r_i+1} \cdots b_1 b_2^{s_i}$$

for each i between 1 and n,

$$a_i^{-1} b_j a_i b_1 b_2^{p_{ij}} b_1 b_2^{p_{ij}+1} \cdots b_1 b_2^{q_{ij}}$$

for i between 1 and m and $j = 1, 2$, and

$$a_i b_j a_i^{-1} b_1 b_2^{u_{ij}} b_1 b_2^{u_{ij}+1} \cdots b_1 b_2^{v_{ij}}$$

for i between 1 and m and $j = 1, 2$, and where for all of these relations the exponents depend on λ.

Observe that for any $\lambda > 0$ one can choose those exponents such that $r_i < s_i$, $p_{ij} < q_{ij}$ and $u_{ij} < v_{ij}$, such that $[r_i, s_i], [p_{ij}, q_{ij}], [u_{ij}, v_{ij}] \subset \mathbb{Z}$ are disjoint, and such that G satisfies $C'(\lambda)$.

The group Q can be made a quotient of G by sending $a_i \mapsto a_i$ and $b_j \mapsto 1$ for $j = 1, 2$. Observe that the second and third collections of relations implies that the subgroup $\langle b_1, b_2 \rangle$ of G is normal in G. So, $K \cong \langle b_1, b_2 \rangle$ is the finitely generated and is the kernel of the quotient map $G \to Q$. □

In [24], Haglund and Wise proved a significant strengthening of the Rips construction, showing that the group G can be taken to effectively embed in a right-angled Artin group. We will not reproduce a proof here.

6.3. **Unsolvability of the isomorphism problem for finitely presented subgroups.** Given the facts we have gathered above, we can now prove Theorem 6.1.

Proof of Theorem 6.1. Let F be a nonabelian free group and let G be a group guaranteed by the Haglund–Wise version of the Rips construction with $Q = F$. We have that G effectively embeds in a right-angled Artin group $A(\Gamma)$. By Theorem 6.2 we have that the isomorphism problem for finitely presented subgroups of $A(\Gamma) \times A(\Gamma) \times F$, which itself is a right-angled Artin group, is unsolvable. It follows that if $A(\Gamma) \times A(\Gamma) \times F$ embeds in $\mathrm{Mod}(S)$ then the isomorphism problem for finitely presented subgroups of $\mathrm{Mod}(S)$ is unsolvable. □

7. Right-angled Artin subgroups of right-angled Artin groups

We have developed a stable, which is to say "up to powers", embedding theory for right-angled Artin groups into mapping class groups. Next, we would like to use that theory in order to develop an analogous theory for right-angled Artin groups themselves. The main reference for this section will be [30].

7.1. **The extension graph.** In the theory of embedding right-angled Artin groups into mapping class groups, the curve graph played a central role in controlling the isomorphism types of right-angled Artin groups which embed in a given mapping class group. The analogous object in the theory of right-angled Artin groups is called the *extension graph*. The extension graph of a graph Γ is denoted by Γ^e.

The vertices of the extension graph are given by $\{v^g \mid v \in V(\Gamma),\, g \in A(\Gamma)\}$. A vertex conjugate v^g can be thought of naturally as an element of $A(\Gamma)$. The edges of Γ^e are given by commutation in $A(\Gamma)$. Thus, $v_1^{g_1}$ and $v_2^{g_2}$ are adjacent if and only if they commute in $A(\Gamma)$.

Proposition 7.1. *The vertices $v_1^{g_1}$ and $v_2^{g_2}$ span an edge of Γ^e if and only if the pair $(v_1^{g_1}, v_2^{g_2})$ is conjugate (via the diagonal action of $A(\Gamma)$) to an edge of Γ.*

Proof. Exercise. □

The reader may already note a bit of a resemblance between the curve graph of a surface and the extension graph of a graph. The mapping class group acts on the curve graph with a finite quotient. Indeed, there is one vertex corresponding to the single mapping class group orbit of nonseparating curves and one vertex for each type of separating curve, where the type is determined by the homeomorphism types of the complementary surfaces. Also, if we identify sccs on a surface with the Dehn twists about them, the curve graph just becomes the commutation graph of the set of Dehn twists in the mapping class group.

A few elementary observations are in order. Note that $A(\Gamma)$ acts on Γ^e in a natural way, and the quotient of Γ^e by this action is just Γ. Observe that $\Gamma = \Gamma^e$ if and only if Γ is complete. We will note more properties of Γ^e below.

There is another construction of the extension graph which is purely combinatorial and does not require recourse to the right-angled Artin group. Let Γ be a finite simplicial graph and let $v \in V(\Gamma)$ be a vertex. Consider the graph
$$\Gamma_v = \Gamma \cup_{st(v)} \Gamma,$$
which is to say the graph built out of two copies of Γ glued together along the star of v by the identity map. The graph Γ_v embeds in Γ^e. Indeed, we can think of one copy of Γ as Γ itself, and the other copy of Γ as Γ^v, the result of conjugating the vertex generators of Γ by v. The conjugation action of v fixes exactly the vertices in the star of v. Thus, the vertices in the star of v are equal in Γ and in Γ^v. The operation $\Gamma \mapsto \Gamma_v$ is called *doubling along the star of v*.

Lemma 7.2. *Let*

$$\Gamma = \Gamma_0 \subset \Gamma_1 \subset \cdots$$

be a sequence of graphs for which:

(1) *For all $i \geq 1$, the graph Γ_i is obtained from Γ_{i-1} by doubling along the star of a vertex of Γ_{i-1}.*
(2) *For each i and each vertex $v \in \Gamma_i$, there exists an infinite sequence of integers*

$$i \leq k_1 < k_2 < \cdots$$

such that Γ_{k_j+1} is obtained from Γ_{k_j} by doubling along the star of the image of v in Γ_{k_j}.

Then there exists an isomorphism of graphs

$$\Gamma^e \cong \bigcup_i \Gamma_i.$$

Proof. Exercise. □

7.2. **Right-angled Artin subgroups of right-angled Artin groups.** We would now like to formulate and prove a result which is analogous to Theorem 5.6 for right-angled Artin groups.

Theorem 7.3. *Let Γ and Λ be finite simplicial graphs.*

(1) *Let $\iota : \Lambda \to \Gamma^e$ be an injective map of graphs which preserves adjacency and non–adjacency. Then for all $N \gg 0$, the map*

$$\iota_{*,N} : A(\Lambda) \to A(\Gamma)$$

defined by $v \mapsto \iota(v)^N$ is injective.
(2) *Suppose there exists an injective map $A(\Lambda) \to A(\Gamma)$. Then Λ occurs as an induced subgraph of $(\Gamma^e)_k$.*

When Γ is triangle–free, Theorem 7.3 can be made into a complete characterization of right-angled Artin subgroups of $A(\Gamma)$:

Theorem 7.4. *Let Γ be triangle free. There exists an injective map $A(\Lambda) \to A(\Gamma)$ if and only if Λ occurs as an induced subgraph of Γ^e.*

Armed with Theorem 5.10, we can give an efficient proof of Theorem 7.3, part (1):

Proof of Theorem 7.3, part (1). Let S be a surface which is large enough to accommodate Γ as a subgraph of the curve graph $\mathcal{C}(S)$ of S. By Theorem 7.3, sufficiently high powers of Dehn twists about curves which span a copy of Γ in $\mathcal{C}(S)$ will generate a copy of $A(\Gamma)$ in $\mathrm{Mod}(S)$. Fixing such a configuration of curves, we let $N \neq 0$ be a power which works. Let $v \in V(\Gamma)$, which we identify with a twisting curve γ_v. The vertex generator v in $A(\Gamma)$ corresponds to the mapping class $T_{\gamma_v}^N$.

The graph $\Gamma \cup_{st(v)} \Gamma$ can be realized in $\mathcal{C}(S)$ by taking the original twisting curves $\{\gamma_i\}$ which span Γ and adding the curves $\{T_{\gamma_v}^N(\gamma_i)\}$. It is clear that the resulting curves span a copy of $\Gamma \cup_{st(v)} \Gamma$ in $\mathcal{C}(S)$. Replacing N by a higher multiple if necessary, the resulting powers of twists generate a copy of $A(\Gamma \cup_{st(v)} \Gamma)$ inside of the original copy of $A(\Gamma)$ in $\mathrm{Mod}(S)$. In view of Lemma 7.2 and the fact that Λ is a finite graph, the conclusion of the result follows. □

One could have proceeded without appealing to Lemma 7.2 by noting that if $\Lambda \subset \Gamma^e$ then Λ lies in a union of finitely many conjugates of Γ, which can be obtained by applying finitely many mapping classes in $\{T_{\gamma_i}^N\}$ to the original collection of curves $\{\gamma_i\}$ and then appealing to Theorem 5.10 to get the desired conclusion.

Before we give a proof of Theorem 7.3, part (2), we need to develop some more background. Let $1 \neq w \in A(\Gamma)$ be cyclically reduced. We consider the support $\Gamma_w \subset \Gamma$ of w. If Γ_w does not split as a nontrivial join then w is called a *pure factor*. Otherwise, we can decompose $\Gamma_w = \Gamma_1 * \Gamma_2$, where both Γ_1 and Γ_2 are nonempty. We write p_1 and p_2 for the two projection maps

$$p_i : A(\Gamma_1) \times A(\Gamma_2) \to A(\Gamma_i).$$

Then we can write $w = p_1(w)p_2(w)$, where $p_i(w)$ both have length strictly less than w. Observe that $p_1(w)$ commutes with $p_2(w)$. By analyzing the supports of $p_i(w)$, we either determine that $p_i(w)$ is a pure factor or we can decompose $p_i(w)$ further as a product of two shorter words which commute with each other. It can be shown that any cyclically reduced $1 \neq w \in A(\Gamma)$ can be written uniquely (up to permutation of the factors) as a product

$$w = w_1 \cdots w_k,$$

where $[w_i, w_j] = 1$ for each i and j and where each w_i is a pure factor.

We say that a collection of conjugates of pure factors $\{w_1, \ldots, w_k\}$ is *irredundant* if $\langle w_i, w_j \rangle$ is not cyclic for each $i \neq j$. The following result is the key observation in the proof of the second part of Theorem 7.3:

Lemma 7.5. *Let $\{w_1, \ldots, w_k\} \subset A(\Gamma)$ be an irredundant collection of conjugates of pure factors with commutation graph Λ. Then Λ embeds as a subgraph of Γ^e.*

Proof. Let $X = \{\gamma_1, \ldots, \gamma_m\}$ be a collection of curves on a surface S which span a copy of Γ in $\mathcal{C}(S)$. We will write $T_1, \ldots, T_m$ for powers of Dehn twists about $\{\gamma_1, \ldots, \gamma_m\}$ which generate a copy of $A(\Gamma)$.

Let $X_i \subset X$ be the support of the cyclic reduction of w_i, viewed as a product of powers of Dehn twists. The element w_i is cyclically reduced by an element $g_i \in A(\Gamma)$, which is identified with a mapping class $\phi_i \in \mathrm{Mod}(S)$. Since w_i is a conjugate of a pure factor, we have that X_i fills a connected subsurface $S_i \subset S$, so that $\phi_i(S_i)$ is the connected subsurface filled by $\phi_i(X_i)$.

Observe that the co–incidence graph of the subsurfaces $\{\phi_1(S_1), \ldots, \phi_k(S_k)\}$ is equal to Λ. The verification of this fact is a useful exercise for the reader. Since $\phi_i(X_i)$ fills $\phi_i(S_i)$, we can build a pseudo-Anosov mapping class ψ_i supported on $\phi_i(S_i)$ which is a product of the elements $\{T_1, \ldots, T_m\}$ (and is therefore identified with an element of $A(\Gamma)$) such that $\langle \psi_i, \psi_j \rangle$ is not cyclic for $i \neq j$. If S_i is annular than ψ_i is just the Dehn twist about the core curve of $\phi_i(S_i)$. Then, choosing an element x_i of $\phi_i(X_i)$, we let $c_{i,N} = \psi_i^N(x_i)$. If N is sufficiently large then the co–incidence graph of $\{c_{1,N}, \ldots, c_{k,N}\}$ is equal to Λ. Since ϕ_i and ψ_i are identified with elements of $A(\Gamma)$, we have that for each N and each i, a nonzero power of a twist about $c_{i,N}$ is identified with a conjugate of a vertex of $A(\Gamma)$. It follows that Λ embeds in Γ^e. □

In general, we might have an embedding of right-angled Artin groups $A(\Lambda) \to A(\Gamma)$ where vertices of Λ are not sent to conjugates of pure factors, but rather to products of commuting conjugates of pure factors. We have already developed the necessary tools for resolving this difficulty in our proof of Theorem 5.6, part (2).

Proof of Theorem 7.3, part (2). Let $\phi : A(\Lambda) \to A(\Gamma)$ be an injective map of right-angled Artin groups. A vertex $\lambda_i \in V(\Lambda)$ gets sent under ϕ to a product $w_1^i \cdots w_{k(i)}^i$ of conjugates of pure factors. By Lemma 7.5, the commutation graph of $\{w_j^i\}$, where the index i ranges over the vertices of Λ and the subscripts range over the pure factors of $\phi(\lambda_i)$, embeds in Γ^e, perhaps after removing any redundancies that may have been introduced.

Let X be the commutation graph of $\{w_j^i\}$. We have an injective map

$$A(\Lambda) \to A(X) \to A(\Gamma)$$

whose composition is ϕ, and which sends each vertex of Λ to a product of vertices supported in a clique of X. The map $A(X) \to A(\Gamma)$ is given by sending a vertex x to the pure factor w_x. Since the composition is injective, the first map $A(\Lambda) \to A(X)$ is injective. Call this map ι.

We claim that Λ embeds in X_k, so that then Λ embeds in $(\Gamma^e)_k$. Indeed, choose an embedding $\iota : A(\Lambda) \to A(X)$ such that $supp(\iota(\lambda_i))$ is a clique for each i and such that

$$\sum_i |supp(\iota(\lambda_i))|$$

is minimal. Then the supports of each $\iota(\lambda_i)$ are all distinct.

Suppose not, and write $\iota(\lambda) = x_1^{p_1} \cdots x_k^{p_k}$ and $\iota(\lambda') = x_1^{q_1} \cdots x_k^{q_k}$ such that $k > 0$,

$$\prod_i p_i \neq 0 \neq \prod_i q_i$$

and $\{x_1, \ldots, x_k\}$ span a clique in X. Note that for each vertex u of Λ adjacent to λ, we have $[\iota(\lambda), \iota(u)] = 1$ and so,

$$supp(\iota(\lambda')) \cup supp(\iota(u)) = supp(\iota(\lambda)) \cup supp(\iota(u))$$

spans a clique; this implies that $[\iota(\lambda'), \iota(u)] = 1$ and λ' is adjacent to u. Hence we have an automorphism $\xi : A(\Lambda) \to A(\Lambda)$ defined by $\xi(\lambda) = \lambda\lambda'^{-1}$ and $\xi(y) = y$ for each $y \in V(\Lambda) \setminus \{\lambda\}$. Let $\eta : A(\Lambda) \to A(\Lambda)$ be an embedding that maps λ to λ^{q_1} while fixing all the other vertices. Then

$$\iota' = \iota \circ \eta \circ \xi^{p_1} : A(\Lambda) \to A(X)$$

is an embedding such that

$$supp(\iota'(\lambda)) \subseteq \{x_2, \ldots, x_k\} \subset supp(\iota(\lambda))$$

and that $\iota'(y) = \iota(y)$ for $y \subseteq V(\Lambda) \setminus \{\lambda\}$. This contradicts the minimality of supports assumption, whence Λ embeds in X_k. □

Theorem 7.3 has many applications. The first corollary which we shall establish recovers a result originally due to Kambites in [26] to which we alluded at the beginning of the course:

Corollary 7.6. *Suppose there exists an injective map $F_2 \times F_2 \to A(\Gamma)$. Then Γ contains an induced square.*

Proof. Label the vertices of a square C_4 cyclically by $\{a, b, c, d\}$, and let $\{X_a, X_b, X_c, X_d\}$ be cliques of Γ_e which realize the embedding $C_4 \to (\Gamma^e)_k$. Choose two pairs of vertices (s, t) and (u, v) in $X_a \times X_c$ and $X_b \times X_d$ which do not span edges in Γ^e. We claim that the graph spanned by $\{s, t, u, v\}$ is a square in Γ^e. Observe that any choice of u and v is adjacent to each choice of s and t. Furthermore, s is distinct from both u and v since the latter two are adjacent to t and s is not. Similarly, t is distinct from both u and v. Thus, these vertices span a square in Γ^e.

We now claim that if Γ^e contains a square then Γ contains a square. This follows more or less from the centralizer theorem. The details are left as an exercise for the reader. □

References

[1] Ian Agol. Criteria for virtual fibering. *J. Topol.*, 1, 269–284, 2008.
[2] A. Baudisch, Subgroups of semifree groups. *Acta Math. Acad. Sci. Hungar.* 38 (1981), no. 1–4, 19–28.
[3] Jason Behrstock and Ruth Charney. Divergence and quasi–morphisms of right-angled Artin groups. To appear in *Math. Annalen.*
[4] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.* 129, 123–139, 1997.
[5] Mladen Bestvina, Kenneth Bromberg and Koji Fujiwara. The asymptotic dimension of mapping class groups is finite. arXiv:1006.1939v2 [math.GR].
[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] Noel Brady. Dehn functions of subgroups of CAT(0) groups. In *The Geometry of the Word Problem for Finitely Generated Groups*. Birkhäuser, 2007.
[8] Martin R. Bridson and Charles F. Miller III. Recognition of subgroups of direct products of hyperbolic groups, *Proc. Amer. Math. Soc.* 132, 59–65, 2004.

[9] Martin R. Bridson. On the subgroups of right-angled Artin groups and mapping class groups. arXiv:1205.5416 [math.GR].

[10] Kenneth S. Brown. *Cohomology of groups.* Graduate Texts in Mathematics, no. 87, Springer, New York, 1982.

[11] M. Casals-Ruiz, A. Duncan, I. Kazachkov. Two examples of partially commutative subgroups of partially commutative groups. In preparation.

[12] Matt Clay, Christopher J. Leininger, and Johanna Mangahas. The geometry of right-angled Artin subgroups of mapping class groups. To appear in *Groups Geom. Dyn.*

[13] John Crisp and Benson Farb. The prevalence of surface groups in mapping class groups. Preprint.

[14] 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.

[15] John Crisp and Bert Wiest. Quasi-isometrically embedded subgroups of braid and diffeomorphism groups. *Trans. Amer. Math. Soc.* 359, no. 11, 5485–5503, 2007.

[16] Michael W. Davis and Tadeusz Januszkiewicz. Right-angled Artin groups are commensurable with right-angled Coxeter groups. *Journal of Pure and Applied Algebra* 153, 229–235, 2000.

[17] Reinhard Diestel. *Graph Theory.* Graduate Texts in Mathematics 173, Springer, 2010.

[18] Carl Droms. Isomorphisms of graph groups. *Proc. Amer. Math. Soc.* 100, no. 3, 407–408, 1987.

[19] Carl Droms. Subgroups of graph groups. *J. Algebra* 110, 519–522, 1987.

[20] Carl Droms, Brigitte Servatius, and Herman Servatius. Surface subgroups of graph groups. *Proc. Amer. Math. Soc.*, 106(3):573–578, 1989.

[21] Benson Farb and Dan Margalit. *A Primer on the mapping class group. Princeton Mathematical Series, 49. Princeton University Press, Princeton, NJ*, 2012.

[22] A. Fathi and F. Laudenbach and V. Poénaru. *Travaux de Thurston sur les Surfaces. Soc. Math. de France, Paris*, Astérisque 66–67, 1979.

[23] M. Gromov. *Hyperbolic Groups*, "Essays in Group Theory" (G. M. Gersten, ed.), MSRI Publ. 8, pp. 75–263, 1987.

[24] Frédéric Haglund and Dani Wise. Special cube complexes. *Geom. Funct. Anal.* 17, no. 5, 1551–1620, 2008.

[25] J. Howie. On the asphericity of the ribbon disc complement. *Trans. Amer. Math. Soc.* 289, no.1, 281–302, 1985.

[26] Mark Kambites. On commuting elements and embeddings of graph groups and monoids. *Proc. Edinb. Math. Soc.* (2) 52, no. 1, 155–170, 2009.

[27] Michael Kapovich and Bernhard Leeb. Actions of discrete groups on nonpositively curved spaces. *Math. Ann.* 306, no. 2, 341–352, 1996.

[28] Ki Hang Kim, L. Makar–Limanov, Joseph Neggers, Fred W. Roush. Graph algebras. *J. Algebra* 64, no. 1, 46–51, 1980.

[29] Felix Klein. Neue Beiträge zur Riemann'schen Functionentheorie. *Math. Ann.* 21 (1883), no. 2, 141–218.

[30] Sang-hyun Kim and Thomas Koberda. Embedability of right-angled Artin groups. To appear in *Geom. Topol.*

[31] Sang-hyun Kim and Thomas Koberda. An obstruction to embedding right-angled Artin groups in mapping class groups. To appear in *Int. Math. Res. Notices.*

[32] Thomas Koberda. Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups. To appear in *Geom. Funct. Anal.*

[33] Thomas Koberda. Ping–pong lemmas with applications to geometry and topology. To appear in *IMS Lecture Notes*, Singapore, 2012.

[34] 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.

[35] Christopher J. Leininger. One–ended subgroups of right-angled Artin groups. Preprint.

[36] R.C. Penner and J.L. Harer. *Combinatorics of train tracks.* Annals of Mathematics Studies, no. 125, *Princeton University Press*, 1992.

[37] C.F. Miller III, *On group–theoretic decision problems and their classification*, Annals of Mathematics Studies, No. 68, Princeton University Press, 1971.

[38] E.S. Rapaport. Note on Nielsen transformations. *Proc. Amer. Math. Soc.*, 10, 229–235, 1959.

[39] E. Rips. Subgroups of small cancellation groups. *Bull. London Math Soc.* 14, 45–47, 1982.

[40] Lucas Sabalka. On rigidity and the isomorphism problem for tree braid groups. *Groups Geom. Dyn.* 3, no. 3, 469–523, 2009.

[41] H. Servatius. Automorphisms of graph groups. *J. Algebra* 126, no. 1, 34–60, 1989.

[42] John Stallings. Group theory and three-dimensional manifolds. A James K. Whittemore Lecture in Mathematics given at Yale University, 1969. Yale Mathematical Monographs, 4. *Yale University Press, New Haven, Conn.–London*, 1971.

Department of Mathematics, Yale University, P.O. Box 208283, New Haven, CT 06520-8283
E-mail address: thomas.koberda@gmail.com