Symplectic structures on right-angled Artin groups: Between the mapping class group and the symplectic group

MATTHEW B DAY

We define a family of groups that include the mapping class group of a genus g surface with one boundary component and the integral symplectic group $\text{Sp}(2g, \mathbb{Z})$. We then prove that these groups are finitely generated. These groups, which we call *mapping class groups over graphs*, are indexed over labeled simplicial graphs with $2g$ vertices. The mapping class group over the graph Γ is defined to be a subgroup of the automorphism group of the right-angled Artin group A_Γ of Γ. We also prove that the kernel of $\text{Aut}\, A_\Gamma \to \text{Aut}\, H_1(A_\Gamma)$ is finitely generated, generalizing a theorem of Magnus.

1 Introduction

1.1 Background

Let Γ be a graph on n vertices, with vertex set X and adjacency relation denoted by $\text{e}(-, -)$. Let A_Γ denote the *right-angled Artin group of* Γ, defined by

$$A_\Gamma := \langle X \mid R_\Gamma \rangle$$

where the relations are $R_\Gamma = \{[x, y] \mid x, y \in X \text{ and } \text{e}(x, y)\}$. As we vary Γ, the group A_Γ interpolates between the free group F_n (if Γ is edgeless) and the free abelian group $\mathbb{Z}^n$ (if Γ is complete). Similarly, as we vary Γ, the automorphism group $\text{Aut}\, A_\Gamma$ interpolates between $\text{Aut}\, F_n$ and the integral general linear group $\text{GL}(n, \mathbb{Z})$.

Both mapping class groups and symplectic groups can be expressed as stabilizer subgroups of automorphism groups. Consider the free group F_{2g} with free generators $a_1, \ldots, a_g, b_1, \ldots, b_g$. The stabilizer in $\text{Aut}\, F_{2g}$ of the element $[a_1, b_1] \cdots [a_g, b_g]$ is a subgroup isomorphic to the mapping class group of a genus g surface with a single boundary component. This is a version, due to Zieschang [20], of the classical Dehn–Nielsen–Baer Theorem (see Farb–Margalit [7, Chapter 3]). At the other extreme, the integral symplectic group $\text{Sp}(2g, \mathbb{Z})$ is the stabilizer in $\text{GL}(2g, \mathbb{Z})$ of the standard symplectic form on $\mathbb{Z}^{2g}$. In this paper, we define a structure on a right-angled Artin

group that interpolates between a surface relator on a free group and a symplectic form, so that the stabilizer in $\operatorname{Aut} A_\Gamma$ of this structure interpolates between mapping class groups and integral symplectic groups. This answers a question due to Benson Farb.

This paper is a continuation of the author's previous paper [6], and we will freely use notation, terminology and results from that paper.

1.2 Symplectic structures on right-angled Artin groups

Let $A'_\Gamma = [A_\Gamma, A_\Gamma]$ be the commutator subgroup of A_Γ. Let $H_\Gamma = A_\Gamma/A'_\Gamma$ be the abelianization of A_Γ, which is the free abelian group $\langle \{[x] \mid x \in X\} \rangle$. As usual, the alternating square $\Lambda^2 H_\Gamma$ of H_Γ is the free abelian group generated by the wedge products $[x] \wedge [y]$ for $x \neq y \in X$ (where $[x] \wedge [y] = -[y] \wedge [x]$). The symbol $[x] \wedge [y]$ is bilinear, so the action of $\operatorname{Aut} A_\Gamma$ on H_Γ induces a diagonal action on $\Lambda^2 H_\Gamma$.

A *standard alternating form* is an element of $\Lambda^2 H_\Gamma$ of the form $[a_1] \wedge [b_1] + \cdots + [a_k] \wedge [b_k]$, where $a_i, b_i \in X^{\pm 1}$ and the $\{a_i, b_i\}_i$ are pairwise distinct and not equal to each other's inverses. A *surface relator* is an element (possibly trivial) of A'_Γ of the form $[a_1, b_1] \cdots [a_k, b_k]$, where $a_i, b_i \in X^{\pm 1}$, and the $\{a_i, b_i\}$ are pairwise distinct and not equal to each other's inverses.

Definition 1.1 Suppose Γ has $2g$ vertices. A pair $(w, Q) \in A_\Gamma \times (\Lambda^2 H_\Gamma)$ is a *symplectic structure* for the right-angled Artin group A_Γ if there is some labeling of $X^{\pm 1}$ as $a_1^{\pm 1}, \ldots, a_g^{\pm 1}, b_1^{\pm 1}, \ldots, b_g^{\pm 1}$ and some k with $0 \leq k \leq g$ satisfying the following conditions:

- For each i with $1 \leq i \leq k$, we have $[a_i, b_i] \neq 1$.

- For each i with $k + 1 \leq i \leq g$, we have $[a_i, b_i] = 1$.

- w is the surface relator

$$w = [a_1, b_1] \cdots [a_k, b_k].$$

- Q is the standard alternating form

$$Q = \sum_{i=k+1}^{g} [a_i] \wedge [b_i].$$

The *mapping class group over* Γ with respect to a symplectic structure (w, Q), written $\operatorname{Mod}(\Gamma, w, Q)$, is the intersection of the stabilizers of w and Q in $\operatorname{Aut} A_\Gamma$:

$$\operatorname{Mod}(\Gamma, w, Q) := (\operatorname{Aut} A_\Gamma)_{(w,Q)} = (\operatorname{Aut} A_\Gamma)_w \cap (\operatorname{Aut} A_\Gamma)_Q$$

Remark 1.2 Consider the subgroups V and $V^\perp$ defined by:

$$V = \langle \{[x] \wedge [y] \mid x, y \in X \text{ and } [x, y] = 1\} \rangle$$

and
$$V^\perp = \langle \{[x] \wedge [y] \mid x, y \in X \text{ and } [x, y] \neq 1\} \rangle$$

As an Aut A_Γ–module, $\Lambda^2 H_\Gamma$ decomposes as $\Lambda^2 H_\Gamma = V \oplus V^\perp$. This phenomenon is somewhat peculiar to right-angled Artin groups.

Let $A_\Gamma^{(2)} = [A_\Gamma', A_\Gamma]$. The map $\Lambda^2 H_\Gamma \to A_\Gamma'/A_\Gamma^{(2)}$ given by $[a] \wedge [b] \mapsto [a, b] \cdot A_\Gamma^{(2)}$ for $a, b \in X$ is a surjective homomorphism (it follows from the Witt–Hall identities that this map is well defined; see Serre [17, Proposition I.II.1.1]). Then V is clearly the kernel of this map. This also tells us that $A_\Gamma'/A_\Gamma^{(2)} \cong V^\perp$. The composition of this isomorphism with the inclusion $V^\perp \hookrightarrow \Lambda^2 H_\Gamma$ gives a map $f \colon A_\Gamma' \to \Lambda^2 H_\Gamma$. Note that if $[a, b] \neq 1$ for $a, b \in X$, then $f([a, b]) = [a] \wedge [b]$. This map f is not usually Aut A_Γ–equivariant because $V^\perp$ is not usually invariant under the action of Aut A_Γ.

If (w, Q) is a symplectic structure on A_Γ, then $Q \in V$, $f(w) \in V^\perp$ and $Q + f(w)$ is a standard symplectic form on H_Γ. It turns out that Aut A_Γ does not usually preserve $Q + f(w)$.

Example 1.3 Suppose Γ is the edgeless graph on $2g$ vertices. Then (w, Q) is a symplectic structure if and only if $Q = 0$ and w is a surface relator of length $4g$. In this case $\mathrm{Mod}(\Gamma, w, 0) \cong \mathrm{Mod}_{g,1}$.

Example 1.4 At the other extreme, if Γ is the complete graph on $2g$ vertices, then (w, Q) is a symplectic structure if and only if $w = 1$ and $Q \in V = \Lambda^2 H_\Gamma$ is a symplectic form. In this case $\mathrm{Mod}(\Gamma, 1, Q) \cong \mathrm{Sp}(2g, \mathbb{Z})$.

The methods of this paper make it possible to explore more difficult examples such as the following, but for brevity we give the following examples without proving the assertions we make about them. We develop an example more thoroughly in Section 4.1.

Example 1.5 Suppose Γ_1 is the complete graph on $2k_1$ vertices, Γ_2 is the edgeless graph on $2k_2$ vertices, and Γ is the graph-theoretic join of Γ_1 and Γ_2. Then a symplectic structure on each of A_{Γ_1} and A_{Γ_2} will induce a symplectic structure (w, Q) on A_Γ. In this case, we have:

$$\mathrm{Mod}(\Gamma, w, Q) \cong \left((\mathrm{Sp}_{2k_1}(\mathbb{Z}) \times \mathrm{Mod}_{k_2,1}) \ltimes \prod_{x \in \Gamma_2} A_{\Gamma_1} \right) \times \mathbb{Z}$$

The inclusions of $\mathrm{Sp}_{2k_1}(\mathbb{Z})$ and $\mathrm{Mod}_{k_2,1}$ into $\mathrm{Mod}(\Gamma, w, Q)$ are the obvious ones, the copy of $\mathbb{Z}$ is given by conjugation by w, and the copies of A_{Γ_1} are given by $x \mapsto xu$ for x a generator in A_{Γ_2} and $u \in A_{\Gamma_1}$.

Example 1.6 If instead of the join, we take Γ to be the disjoint union of the graphs Γ_1 and Γ_2 from Example 1.5, then we still get an induced symplectic structure (w, Q), but a different group $\mathrm{Mod}(\Gamma, w, Q)$. Any automorphism in $\mathrm{Aut}\, A_\Gamma$ that conjugates all of the generators of A_{Γ_1} by one of the generators of A_{Γ_2} (and fixes the generators of A_{Γ_2}) preserves the symplectic structure (w, Q). Although it is not obvious, it turns out that:

$$\mathrm{Mod}(\Gamma, w, Q) \cong \mathrm{Sp}_{2k_1}(\mathbb{Z}) \times (\mathrm{Mod}_{k_2,1} \ltimes A_{\Gamma_2})$$

In general, the groups $\mathrm{Mod}(\Gamma, w, Q)$ and $\mathrm{Mod}(\Gamma', w', Q')$ tend to look very different for slightly different graphs Γ and Γ'. Even if $\Gamma = \Gamma'$, if (w, Q) is different from (w', Q'), the resulting groups may be different.

Example 1.7 If Γ is the disjoint union of Γ_1 and Γ_2 as in Example 1.6, but with a single edge added between a vertex of Γ_1 and a vertex of Γ_2, then the respective inclusions of Γ_1 and Γ_2 into Γ do not induce maps of $\mathrm{Aut}\, A_{\Gamma_1}$ or $\mathrm{Aut}\, A_{\Gamma_2}$ to $\mathrm{Aut}\, A_\Gamma$. Then neither $\mathrm{Mod}_{k_2,1}$ nor $\mathrm{Sp}(2k_1, \mathbb{Z})$ include in $\mathrm{Mod}(\Gamma, w, Q)$ in the obvious way. However, $\mathrm{Mod}_{k_2-1,1}$ and $\mathrm{Sp}(2k_1 - 2, \mathbb{Z})$ both include into $\mathrm{Mod}(\Gamma, w, Q)$, so the group is nontrivial. It takes some work to get a more complete picture of this group.

1.3 Statement of results

Theorem A *For any graph Γ with an even number of vertices and any symplectic structure (w, Q) on A_Γ, the group $\mathrm{Mod}(\Gamma, w, Q)$ is finitely generated.*

This is strong evidence that our definition for $\mathrm{Mod}(\Gamma, w, Q)$ from Definition 1.1 is a good one. We also considered an alternate definition for a symplectic structure: a pair (w, Q) where w is a surface relator and $Q \in \Lambda^2 H_\Gamma$ is a standard symplectic form, such that w and Q project to the same element in $A'_\Gamma / A_\Gamma^{(2)}$. As we show in Section 4.1, there is an example of a graph Γ where the subgroup of $\mathrm{Aut}\, A_\Gamma$ fixing both a surface relator and a compatible symplectic form on H_Γ is not finitely generated (but of course, Theorem A still holds in this case).

The proof of Theorem A proves both the finite generation of mapping class groups and the integral symplectic groups in special cases. We did not find a single argument that proved both things in the same way, but rather found a single algorithm that reduces to two previously known algorithms in each extreme case. These extremal algorithms are integral symplectic row reduction and the peak reduction algorithm (Whitehead's theorem) for free groups.

We also obtain the following statement, which is of interest in itself, as a corollary to a proposition used in the proof of Theorem A. In the case where A_Γ is a free group,

this theorem restricts to the 1934 theorem of Magnus that $\ker(\operatorname{Aut} F_n \to \operatorname{GL}(n, \mathbb{Z}))$ is finitely generated (see Magnus–Karrass–Solitar [12, Theorem 3.5.N4] or Magnus [11]). Let $\operatorname{IAut} A_\Gamma$ denote the kernel $\ker(\operatorname{Aut} A_\Gamma \to \operatorname{Aut} H_\Gamma)$.

Theorem B *The group* $\operatorname{IAut} A_\Gamma$ *is finitely generated.*

This theorem opens the way for further study of $\operatorname{IAut} A_\Gamma$. An interesting corollary of this theorem is that the preimage in $\operatorname{Aut} A_\Gamma$ of a finitely generated subgroup of $\operatorname{Im}(\operatorname{Aut} A_\Gamma \to \operatorname{Aut} H_\Gamma)$ is a finitely generated group.

2 Background

In this section, we review the notation and the main result from Day [6]. Let L be the set of letters $X \cup X^{-1}$. For $x \in L$, let $\operatorname{v}(x) \in X$, the *vertex of x*, be the unique element of $X \cap \{x, x^{-1}\}$. We will use $\operatorname{st}_L(x)$ and $\operatorname{lk}_L(x)$ as notation for $\operatorname{st}(\operatorname{v}(x)) \cup \operatorname{st}(\operatorname{v}(x))^{-1}$ and $\operatorname{lk}(\operatorname{v}(x)) \cup \operatorname{lk}(\operatorname{v}(x))^{-1}$ respectively. The *support* $\operatorname{supp} w$ of a word or cyclic word w is the subset of X consisting of all generators that appear (or whose inverses appear) in w. There is a reflexive and transitive binary relation on X called the *domination relation*: say $x \geq y$ (x dominates y) if $\operatorname{lk}(y) \subset \operatorname{st}(x)$. Write $x \sim y$ when $x \geq y$ and $y \geq x$; the relation $\sim$ is called the *domination equivalence* relation.

There are four important classes of automorphisms known collectively as the Laurence–Servatius generators: dominated transvections, partial conjugations, inversions and graphic automorphisms. For $x, y \in L$ with $x \geq y$ and $\operatorname{v}(x) \neq \operatorname{v}(y)$, the *dominated transvection* (or simply *transvection*) $\tau_{x,y}$ is the automorphism that sends $y \mapsto yx$ and fixes all generators not equal to $\operatorname{v}(y)$. For $x \in L$ and Y a union of connected components of $\Gamma - \operatorname{st}(\operatorname{v}(x))$, the *partial conjugation* $c_{x,Y}$ is the automorphism that sends $y \mapsto x^{-1}yx$ for $y \in Y$ and fixes all generators not in Y. For $x \in X$, the *inversion*

of x is the automorphism that sends $x \mapsto x^{-1}$ and fixes all other generators. For π an automorphism of the graph Γ, the *graphic automorphism* of π sends $x \mapsto \pi(x)$ for each generator $x \in X$. Servatius defined these automorphisms and conjectured that they generate $\operatorname{Aut} A_\Gamma$ in [18]; Laurence proved that conjecture in [9, Theorem 6.9].

We will use the *Whitehead automorphisms* of $\operatorname{Aut} A_\Gamma$, as defined by the author in [6]. The set of Whitehead automorphisms Ω is the finite set of all automorphisms of the following two types. The *type (1)* Whitehead automorphisms are the finite subgroup of $\operatorname{Aut} A_\Gamma$ generated by the inversions and graphic automorphisms. An automorphism $\alpha \in \operatorname{Aut} A_\Gamma$ is a *type (2)* Whitehead automorphism if there is an element $a \in L$, called the *multiplier* of α, such that for all $x \in X$, we have $\alpha(x) \in \{x, xa, a^{-1}x, a^{-1}xa\}$ (note $\alpha(a) = a$). For $a \in L$ and $A \subset L$ with $a \in A$ and $a^{-1} \notin A$, we use the notation (A, a) to refer to the type (2) Whitehead automorphism that sends $x \in L - \{a, a^{-1}\}$ to x or $a^{-1}x$ if $x \notin A$ and to xa or $a^{-1}xa$ if $x \in A$, if such an automorphism exists. Lemma 2.5 of Day [6] explains when such an automorphism exists.

The following two subsets of Ω are also from Day [6]. The set Ω_ℓ of *long-range* Whitehead automorphisms is the set of all type (1) Whitehead automorphisms together with all type (2) Whitehead automorphisms (A, a) with $A \cap \operatorname{lk}_L(a) = \varnothing$. The set Ω_s of *short-range* Whitehead automorphisms is the set of type (2) Whitehead automorphisms (A, a) with $A \subset \operatorname{st}_L(a)$.

We recall the definition of peak reduction. The length of a conjugacy class in $\operatorname{Aut} A_\Gamma$ is the shortest length of a representative element (with respect to X). We say that a factorization $\alpha = \beta_k \cdots \beta_1$ is peak-reduced with respect to a conjugacy class $[w]$ in $\operatorname{Aut} A_\Gamma$ if for each $i = 1, \ldots k$, we do not have both

$$|\beta_{i+1} \cdots \beta_1([w])| \le |\beta_i \cdots \beta_1([w])|$$

and
$$|\beta_i \cdots \beta_1([w])| \ge |\beta_{i-1} \cdots \beta_1([w])|$$

unless all three lengths are equal. We say that α can be peak-reduced by elements of a set S with respect to $[w]$ if there is a factorization $\alpha = \beta_k \cdots \beta_1$ by elements $\beta_1, \ldots, \beta_k$ that is peak-reduced with respect to $[w]$.

The following theorem is essentially Theorem B of Day [6].

Theorem 2.1 *The set $\Omega_\ell \cup \Omega_s$ is a finite generating set for $\operatorname{Aut} A_\Gamma$ with the following properties:*

(1) *Each $\alpha \in \operatorname{Aut} A_\Gamma$ can be written as $\alpha = \beta\gamma$ for some $\beta \in \langle \Omega_s \rangle$ and some $\gamma \in \langle \Omega_\ell \rangle$.*

(2) *The usual representation $\operatorname{Aut} A_\Gamma \to \operatorname{Aut} H_1(A_\Gamma)$ to the automorphism group of the abelianization $H_1(A_\Gamma)$ of A_Γ restricts to an embedding $\langle \Omega_s \rangle \hookrightarrow \operatorname{Aut} H_1(A_\Gamma)$.*

(3) *Any $\alpha \in \langle \Omega_\ell \rangle$ can be peak-reduced by elements of Ω_ℓ with respect to any conjugacy class $[w]$ in A_Γ.*

We will also make use of the *pure automorphism group* of A_Γ, denote $\mathrm{Aut}^0 A_\Gamma$. The group $\mathrm{Aut}^0 A_\Gamma$ is the subgroup of $\mathrm{Aut}\, A_\Gamma$ generated by the partial conjugations, dominated transvections and inversions. This group appears in Charney–Crisp–Vogtmann [4] and is useful for technical reasons. It is easy to see that $\mathrm{Aut}^0 A_\Gamma$ is normal in $\mathrm{Aut}\, A_\Gamma$ and that $\mathrm{Aut}\, A_\Gamma / \mathrm{Aut}^0 A_\Gamma$ is finite (it is a quotient of $\mathrm{Aut}\,\Gamma$). The group $\mathrm{Aut}^0 A_\Gamma$ contains all those graphic automorphisms that can be realized as a product of transvections and inversions, so if Γ is edgeless or complete, then $\mathrm{Aut}^0 A_\Gamma$ is $\mathrm{Aut}\, A_\Gamma$.

3 Kernels of restrictions of the homology representation

This section is devoted to the proof of Theorem B. We will also prove a proposition that will be used in the proof of Theorem A.

If $x, y, c \in L$ with $x, y \geq c$ and $\mathrm{v}(x)$, $\mathrm{v}(y)$ and $\mathrm{v}(c)$ all distinct, then we write $\tau_{[x,y],c}$ for $[\tau_{x,c}, \tau_{y,c}]$. As the notation suggests, $\tau_{[x,y],c}$ sends $c \mapsto c[x, y]$ and fixes all generators in X not equal to c.

For any subset $Z \subset X$, let $G_Z < \mathrm{Aut}\, A_\Gamma$ be generated by the transvections $\tau_{a,b}$ for $a, b \in Z^{\pm 1}$ with $a \geq b$, and the (total) conjugations of A_Γ. Let $K_Z < \mathrm{Aut}\, A_\Gamma$ be generated by all the $\tau_{[x,y],c}$ and all the partial conjugations $c_{x,\{c\}}$ for $x, y, c \in Z$ with $x, y \geq c$, and the (total) conjugations of A_Γ. Note that for each Z, we have $K_Z < G_Z$. We will refer to a partial conjugation of the form $c_{x,\{c\}}$ as a *one-term partial conjugation*.

Remark 3.1 In fact, K_Z is equal to the subgroup generated by the conjugations and the $\tau_{[x,y],c}$ and $c_{x,\{d\}}$ for $x, y, c \in Z^{\pm 1}$ and $d \in Z$ (with appropriate domination conditions). This is because $c_{x,\{d\}}^{-1} = c_{x^{-1},\{d\}}$, and because $\tau_{[x,y],c}$ with $x, y, c \in Z^{\pm 1}$ can always be expressed as a product of generators of K_Z.

Sublemma 3.2 For any $a, b, c, x \in X$, with $a \geq b$, $a \neq b$, $x \geq c$ and $x \neq c$, the automorphism $\tau_{a,b} c_{x,\{c\}} \tau_{a,b}^{-1}$ is in $K_{\{a,b,c,x\}}$.

For $Y \subset X$ such that $c_{x,Y}$ is a partial conjugation of A_Γ, we have that $\tau_{a,b} c_{x,Y} \tau_{a,b}^{-1}$ is a product of elements of $K_{\{a,b,x\}}$ and partial conjugations of the form $c_{z,Y'}$ where $z \in \{a, x\}$ and $Y' \subset Y \cup \{x, a\}$.

If c_x is conjugation by x, then $\tau_{a,b} c_x \tau_{a,b}^{-1}$ is in $K_{\{a,b,x\}}$.

Proof Suppose that $a, b, x \in X$, $a \geq b$ and $Y \subset X$ such that $c_{x,Y}$ is a partial conjugation. The lemma will follow from several identities of automorphisms, which can be verified by evaluating the automorphisms on X. Note that if $a = x$, then $\tau_{a,b}$ and $c_{x,Y}$ commute. If both $a, b \in Y$, then $\tau_{a,b}$ and $c_{x,Y}$ commute. If $a \in Y$, $b \notin Y$ and $b \neq x$, then the following identity applies:

$$(3\text{-}1) \qquad\qquad \tau_{a,b}c_{x,Y}\tau_{a,b}^{-1} = c_{x,Y}\tau_{[x,a],b}$$

The use of $\tau_{[x,a],b}$ is allowed, since if a and b are in different components of $\Gamma - \text{st}(x)$ and $a \geq b$, then $x \geq b$. If $a \in Y$ and $b = x$, then:

$$(3\text{-}2) \qquad\qquad \tau_{a,x}c_{x,Y}\tau_{a,x}^{-1} = c_{a,(Y-a+x)}c_{x,Y}$$

These terms are allowed since if $a \geq x$ and Y is a union of connected components of $\Gamma - \text{st}(x)$, then $Y - a$ and $Y - a + x$ are both unions of connected components of $\Gamma - \text{st}(a)$. We have covered all the cases where $a \in Y$ or $a = x$, so we assume that $a \notin Y$ and $a \neq x$. If both $a, b \notin Y$ and $x \neq a$ and $x \neq b$, then $\tau_{a,b}$ and $c_{x,Y}$ commute. If $b \in Y$, then:

$$(3\text{-}3) \qquad\qquad \tau_{a,b}c_{x,Y}\tau_{a,b}^{-1} = c_{x,Y}\tau_{[x^{-1},a],b}$$

As in Equation (3-1), the conditions ensure that $\tau_{[x^{-1},a],b}$ is allowed. If $b = x$, then:

$$(3\text{-}4) \qquad\qquad \tau_{a,x}c_{x,Y}\tau_{a,x}^{-1} = c_{x,Y}c_{a,Y}$$

Since $a \geq x$ and $a \notin Y$, we have that Y is a union of connected components of $\Gamma - \text{st}(a)$ and $c_{a,Y}$ is allowed. This proves the second statement in the lemma.

If we have some $c \in X$ with $x \geq c$, then we can take $Y = \{c\}$ and each of the equations from Equation (3-1) through Equation (3-4) applies, proving the first statement in the lemma.

The third statement is obvious since the groups of inner automorphisms is normal in $\text{Aut}\, A_\Gamma$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Sublemma 3.3 *For any $a, b, c, x, y \in X$, with $a \geq b$, $a \neq b$, $x \geq c$, $x \neq c$, $y \geq c$, $y \neq c$ and $x \neq y$, the automorphism $\tau_{a,b}\tau_{[x,y],c}\tau_{a,b}^{-1}$ is in $K_{\{a,b,c,x,y\}}$.*

Proof Note that

$$\tau_{[y,x],c}^{-1} = \tau_{[x,y],c}$$

so we may switch x and y in our enumeration of cases.

If $a, b \notin \{x, y, c\}$, then it follows from Day [6, Proposition 2.10] $\tau_{a,b}$ and $\tau_{[x,y],c}$ commute. If $c = b$ then one can verify by evaluation on X that:

$$\tau_{a,b} \tau_{[x,y],b} \tau_{a,b}^{-1} = c_{a,\{b\}} \tau_{[x,y],b} c_{a,\{b\}}^{-1}$$

This works whether or not $a \in \{x, y\}$. If $c = a$ and $b \notin \{x, y\}$, then it follows from Day [6, Proposition 2.10] and the previous case that:

$$\tau_{a,b} \tau_{[x,y],a} \tau_{a,b}^{-1} = \tau_{[x,y],a} c_{a,\{b\}} \tau_{[y,x],b} c_{a,\{b\}}^{-1}$$

If $c = a$ and $b = x$, then

$$\tau_{a,b} \tau_{[b,y],a} \tau_{a,b}^{-1} = c_{a,\{b\}} c_a^{-1} c_{b,\{a\}} c_b^{-1} c_{y,\{b\}}^{-1} \tau_{[y,a],b} c_{y,\{a\}} \tau_{[y^{-1},b^{-1}],a^{-1}} c_b c_{b,\{a\}}^{-1} c_a c_{a,\{b\}}^{-1}$$

where c_a and c_b denote the (total) conjugations by a and b respectively. Our assumptions dictate that $y \geq a \sim b$, so all the terms in this equation are allowed. This identity can be verified by evaluation on X.

Now we may assume that $c \notin \{a, b\}$. If $b = x$ and $a \neq y$, then the following identity applies:

$$\tau_{a,b} \tau_{[b,y],c} \tau_{a,b}^{-1} = c_{b,\{c\}} \tau_{[a,y],c} c_{b,\{c\}}^{-1} \tau_{[b,y],c}$$

Again, this identity can be verified by evaluation. There are then two remaining cases: $a = x$ and $b \neq y$; and $a = x$ and $b = y$. In both of these cases, it follows from Day [6, Proposition 2.10] that $\tau_{a,b}$ commutes with $\tau_{[x,y],c}$. □

Lemma 3.4 *For any $Z \subset X$, the group K_Z is normal in G_Z.*

Proof If $a, b \in X$ with $a \geq b$, then $\tau_{a^{-1},b} = \tau_{a,b}^{-1}$, and $\tau_{a,b^{-1}} \tau_{a,b} = c_{a,\{b\}}$. This means that G_Z is generated by the generators of K_Z together with the transvections $\tau_{a,b}$ with $a, b \in K$ (in particular, not in K^{-1}). Then the identities from Sublemma 3.2 and Sublemma 3.3 indicate that the conjugate of any generator of K_Z by a generator of G_Z can be expressed as a product of elements of K_Z. □

The proof of the following proposition is a generalization of Magnus's proof that IA_n is finitely generated [11].

Proposition 3.5 *For any $Z \subset X$, we have $K_Z = \ker(G_Z \to \operatorname{Aut} H_\Gamma)$.*

Proof Let $C_1 \cup \cdots \cup C_m = Z$ be the decomposition of Z into domination equivalence classes. Since partial conjugations map to the identity in $\operatorname{Aut} H_\Gamma$, it follows from Day [6, Corollary 3.11] that $\operatorname{Im}(G_Z \to \operatorname{Aut} H_\Gamma)$ has a presentation where the generators are the elementary row operations $E_{a,b} = (\tau_{a,b})_*$ such that $a \geq b$, for $a, b \in Z$, and the relations are as follows:

(1) $[E_{a,b}, E_{c,d}] = 1$ if $b \neq c$ and $a \neq b$.

(2) $[E_{a,b}, E_{b,d}]E_{a,d}^{-1} = 1$ if $a \neq d$.

(3) $(E_{a,b}E_{b,a}^{-1}E_{a,b})^4 = 1$, if $a \sim b$ and $a \neq b$.

(4) $(E_{a,b}E_{b,a}^{-1}E_{a,b})^2(E_{a,b}E_{b,a}^{-1}E_{a,b}E_{b,a})^{-3} = 1$, if $a, b \in C_i$, $a \neq b$ and $|C_i| = 2$.

Consider the lifts of the relations gotten by replacing each of the $E_{a,b}$ with the corresponding $\tau_{a,b}$. We claim that these lifts are all in K_Z. Relation (1) obviously lifts to $\tau_{[a,c],b}$ if $b = d$ and lifts to the trivial element otherwise. Relation (2) lifts to $[\tau_{a,b}, \tau_{b,d}]\tau_{a,d}^{-1}$, which is $\tau_{[b,a],d}$. We know K_Z is normal in G_Z, so we say two elements of G_Z are equal modulo K_Z if their difference is in K_Z. Since $\tau_{a,b^{-1}}\tau_{a,b} = c_{a,\{b\}}$, we know that $\tau_{a,b}^{-1}$ and $\tau_{a,b^{-1}}$ are equal modulo K_Z. Then the lift the element $E_{a,b}E_{b,a}^{-1}E_{a,b}$ is equal, modulo K_Z, to $\tau_{a,b^{-1}}^{-1}\tau_{b,a^{-1}}\tau_{a,b}$, which is equal to the permutation $\sigma_{a,b}$ of order 4 from Equation (R5) of Day [6], according to that equation. So Relation (3) lifts to an element of K_Z. The lift of the element $(E_{a,b}E_{b,a}^{-1}E_{a,b}E_{b,a})^3$ is equal modulo K_Z to $(\sigma_{a,b}\tau_{b,a})^3$. By Equation (R6) of Day [6], $(\sigma_{a,b}\tau_{b,a})^3 = \tau_{a,b}^{-1}\tau_{b,a^{-1}}^{-1}\tau_{a,b^{-1}}\sigma_{a,b}^3$. This is equal modulo K_Z to $\tau_{a,b}^{-1}\tau_{b,a^{-1}}^{-1}\tau_{a,b}^{-1}\sigma_{a,b}^3$, which is σ^2 by Equation (R5) of Day [6]. So Relation (4) lifts to an element of K_Z.

The group K_Z is obviously in $\ker(G_Z \to \operatorname{Aut} H_\Gamma)$. Any element of G_Z can be expressed as a product of inner automorphisms, one-term partial conjugations and lifts $\{\tau_{a,b} \mid a, b \in Z, a \geq b\}$ of the $\{E_{a,b} \mid a, b \in Z, a \geq b\}$. Since these lifts map to the generators of our presentation for $\operatorname{Im}(G_Z \to \operatorname{Aut} H_\Gamma)$ and the inner automorphisms and one-term partial conjugations are in $\ker(G_Z \to \operatorname{Aut} H_\Gamma)$, it follows that any element of $\ker(G_Z \to \operatorname{Aut} H_\Gamma)$ can be written as a product of conjugates of inner automorphisms, one-term partial conjugations and lifts of relators from the presentation. The group K_Z contains all the inner automorphisms, one-term partial conjugations and lifts of the relators. By Lemma 3.4, K_Z is normal in G_Z, so it contains all the conjugates of these elements. So $\ker(G_Z \to \operatorname{Aut} H_\Gamma) < K_Z$, and they are equal.   □

Recall from the introduction that $\operatorname{IAut} A_\Gamma$ denotes the kernel of the homology representation. We will show Theorem B by showing that $\operatorname{IAut} A_\Gamma$ is generated by the generators of K_X, together with the partial conjugations of A_Γ.

Proof of Theorem B Let $\rho \colon \operatorname{Aut} A_\Gamma \to \operatorname{Aut} H_\Gamma$ be the homology representation. As previously noted, $\operatorname{Aut}^0 A_\Gamma$ is normal in $\operatorname{Aut} A_\Gamma$. It is apparent from considering the generators of $\operatorname{Aut} A_\Gamma$ and the definition of $\operatorname{Aut}^0 A_\Gamma$ that ρ induces an isomorphism $\operatorname{Aut} A_\Gamma / \operatorname{Aut}^0 A_\Gamma \cong \rho(\operatorname{Aut} A_\Gamma)/\rho(\operatorname{Aut}^0 A_\Gamma)$. From this we deduce that $\operatorname{IAut} A_\Gamma < \operatorname{Aut}^0 A_\Gamma$.

Let K be the subgroup of $\text{Aut}^0 A_\Gamma$ generated by K_X and the partial conjugations. Note that $\text{Aut}^0 A_\Gamma$ is generated by G_X together with K and the inversion automorphisms.

By Sublemma 3.2, Lemma 3.4 and the fact that inversions normalize K, we know that K is normal in $\text{Aut}^0 A_\Gamma$. So if $\alpha \in \text{Aut}^0 A_\Gamma$, then α can be written as $\alpha = \beta\gamma$ where $\beta \in K$ and γ is a product of elements of G_X and inversions. If we further assume that α is in $\text{IAut}\, A_\Gamma$, then it follows from Proposition 3.5 that γ is in K_X. So $\text{IAut}\, A_\Gamma < K$. Since the reverse inclusion is obvious, it follows that $\text{IAut}\, A_\Gamma = K$ and $\text{IAut}\, A_\Gamma$ is generated by the finite set of the generators of K_X together with the partial conjugations. □

4 Symplectic structures

4.1 A counterexample

As an alternate definition for a symplectic structure on a right-angled Artin group, one can consider a pair $(w, \widetilde{Q})$ where $w \in A'_\Gamma$ is a surface relator, $\widetilde{Q} \in \Lambda^2 H_\Gamma$ is a symplectic form, and w and $\widetilde{Q}$ map to the same element under the respective maps of A'_Γ and $\Lambda^2 H_\Gamma$ to $A'_\Gamma / A_\Gamma^{(2)}$. The group $(\text{Aut}\, A_\Gamma)_{(w, \widetilde{Q})}$ can also be seen as an analogue to a mapping class group or a symplectic group. This differs from Definition 1.1 in that $\widetilde{Q}$ is a symplectic form on all of $\Lambda^2 H_\Gamma$, instead of being an alternating form supported on a subspace.

This alternate definition is attractive because the groups defined in this way have symplectic homology representations, while in general the groups $\text{Mod}(\Gamma, w, Q)$ do not. However, this alternate definition is less attractive because of the following example, which is a group that satisfies the alternate definition and is not finitely generated.
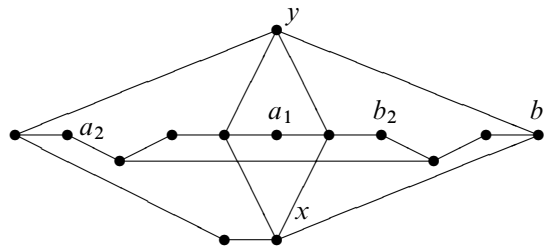

Figure 1: A counterexample to the finite generation of a different group

Example 4.1 Take Γ to be the graph indicated in Figure 1. Let w be the word $[a_1, b_1][a_2, b_2]$, where a_1, b_1, a_2 and b_2 are as indicated. By pairing off the remaining

vertices in adjacent pairs $a_3, b_3, \ldots, a_7, b_7$ (which include the vertices labeled x and y), we get a symplectic form:

$$\tilde{Q} = \sum_{i=1}^{7} [a_i] \wedge [b_i]$$

It is easy to see that $(w, \tilde{Q})$ satisfies the alternate definition.

One can check that the only examples of domination in this Γ are $x \geq a_1$ and $y \geq a_1$, and that the only vertices whose stars separate Γ are x and y, both of which separate it into $\{a_1\}$ and one other component. By inspecting the valences of the vertices, it is apparent that there are no nontrivial automorphisms of the graph Γ.

From Laurence's theorem [9, Theorem 6.9], we can tell that Aut A_Γ is generated by conjugations, inversions and the following four transvections:

$$\{\tau_{x,a_1}, \tau_{y,a_1}, \tau_{x,a_1^{-1}}, \tau_{y,a_1^{-1}}\}.$$

Note that $\tau_{x,a_1}(a_1) = a_1 x$ and $\tau_{y,a_1}(a_1) = a_1 y$, while $\tau_{x,a_1^{-1}}(a_1) = x^{-1} a_1$ and $\tau_{y,a_1^{-1}}(a_1) = y^{-1} a_1$. Let F_2 denote the free group on the generators x and y. If α is in the subgroup generated by these four transvections, then $\alpha(a_1) = u^{-1} a_1 v$ for some $u, v \in F_2$. The map $\alpha \mapsto (u, v)$ is an isomorphism from this subgroup to $F_2 \times F_2$. Then we have

(4-1)   Out $A_\Gamma \cong (\mathbb{Z}/2\mathbb{Z})^{14} \ltimes (F_2 \times F_2)$

where the fourteen generators of order 2 are the inversions and the inversions act on the transvections by the rule of Equation (R6) of Day [6].

If $(u, v) \in F_2 \times F_2$, then the corresponding outer automorphism sends the conjugacy class $[w]$ to the class represented by the graphically reduced cyclic word $u^{-1} a_1 b_1 a_1^{-1} u b_1^{-1} [a_2, b_2]$. The v does not appear because x and y both commute with b_1. It then follows from Equation (4-1) that $(\text{Out } A_\Gamma)_{[w]}$ is the subgroup generated by the images of $\tau_{x,a_1}, \tau_{y,a_1}$ together with the inversions of vertices that do not appear in w. The only inner automorphisms fixing w are conjugation by powers of w. At this point, we can see that

$$(\text{Aut } A_\Gamma)_w \cong ((\mathbb{Z}/2\mathbb{Z})^{10} \ltimes F_2) \times \mathbb{Z}$$

where the copy of F_2 is generated by τ_{x,a_1} and τ_{y,a_1} and the copy of $\mathbb{Z}$ is generated by conjugation by w.

In the subgroup of Aut H_Γ generated by $(\tau_{x,a_1})_*$ and $(\tau_{y,a_1})_*$, it is easy to check that only the trivial element preserves $\tilde{Q}$. Then it follows from Proposition 3.5 (with

$Z = \{x, y, a_1\}$) that the group $(\mathrm{Aut}\, A_\Gamma)_{\tilde{Q}}$ is also finitely generated. We do not see any groups that are not finitely generated until we stabilize both w and $\tilde{Q}$.

As before, if $\alpha \in \langle \tau_{x,a_1}, \tau_{y,a_1} \rangle$, then $\alpha(a_1) = a_1 u$, where $u \in \langle x, y \rangle \cong F_2$ and the map $\alpha \mapsto u$ is an isomorphism. The subgroup of $\langle \tau_{x,a_1}, \tau_{y,a_1} \rangle$ fixing Q is then isomorphic to the kernel of the abelianization map $F_2 \to \mathbb{Z}^2$. The only products of inversions preserving both w and $\tilde{Q}$ are $N_{a_i} N_{b_i}$ for $i = 3, \ldots, 7$ and their products, where N_z denotes the inversion with respect to z. We can then deduce that:

$$(\mathrm{Aut}\, A_\Gamma)_{(w,\tilde{Q})} \cong ((\mathbb{Z}/2\mathbb{Z})^5 \ltimes (\ker(F_2 \to \mathbb{Z}^2))) \times \mathbb{Z}$$

Since $\ker(F_2 \to \mathbb{Z}^2)$ is an infinite rank free group, we have that $(\mathrm{Aut}\, A_\Gamma)_{(w,\tilde{Q})}$ is not finitely generated.

On the other hand, if we take Q to be $\tilde{Q}$ minus the image of w in $\Lambda^2 H_\Gamma$, then

$$\mathrm{Mod}(\Gamma, w, Q) \cong ((\mathbb{Z}/2\mathbb{Z})^5 \ltimes F_2) \times \mathbb{Z},$$

which is finitely generated.

Remark 4.2 This example shows that the image of $\mathrm{Mod}(\Gamma, w, Q)$ under the homology representation $\mathrm{Aut}\, A_\Gamma \to \mathrm{Aut}\, H_\Gamma$ need not be symplectic. Although $\mathrm{Mod}(\Gamma, w, Q)$ fixes w and fixes Q, it doesn't necessarily fix $\tilde{Q} = f(w) + Q$, where $f \colon A'_\Gamma \to \Lambda^2 H_\Gamma$ is as in Remark 1.2. This is because the map f is not usually equivariant.

4.2 Symplectic row reduction with domination

At this point, we assume that $|X| = 2g$ is even. Pick a bijection $()^* \colon L \to L$ such that $(a^*)^* = a^{-1}$ for all $a \in L$, and pick a set of g letters $S = \{a_1, \ldots, a_g\} \subset L$ such that $S^* \cup S$ contains x or x^{-1} for each $x \in X$. Let

$$Q = \sum \left\{ [a_i] \wedge [a_i^*] \mid a_i^* \in \mathrm{lk}_L(a_i) \right\} \in \Lambda^2 H_\Gamma$$

and let w_0 be the concatenation of the words $[a_i, a_i^*]$ for those i for which $a_i^* \notin \mathrm{lk}_L(a_i)$, in increasing order of the index i. Then (w_0, Q) satisfies the definition of a symplectic structure on A_Γ. We will also demand that there is some k, $0 \le k \le g+1$, such that $[a_i, a_i^*] \ne 1$ for $i = 1, \ldots, k$ and $[a_i, a_i^*] = 1$ for $i = k+1, \ldots, g$. In this subsection, we assume $Q \ne 0$. Let $\mathrm{supp}\, Q \subset X$ denote the set of elements $a \in X$ with a appearing in Q.

In this subsection, for a in L, we will also use a to denote the image of a in H_Γ. The images of the elements of X give a basis for H_Γ which we also call X. By declaring X to be orthonormal, we determine an inner product $\langle -, - \rangle \colon H_\Gamma \times H_\Gamma \to \mathbb{Z}$.

For $a, b \in L$ with $a \geq b$, let $E_{a,b} \in \operatorname{Aut} H_\Gamma$ denote the linear transvection (the row operation) mapping

$$b \mapsto b + a$$

and fixing the images of all elements of X different from $\mathrm{v}(a)$. For $a \in L$, let $N_a \in \operatorname{Aut} H_\Gamma$ denote the inversion with respect to a, which maps

$$a \mapsto -a$$

and fixes the images of all elements of X different from $\mathrm{v}(a)$.

A *standard dominated Q–transvection* (or *Q–transvection* for short) is an element of $\operatorname{Aut} H_\Gamma$ of one of the two following forms:

(1) E_{a,a^*}, where $a \in \operatorname{supp} Q$ and $a \geq a^*$.

(2) $E_{a,b} E_{b^*,a^*}^{-1}$ where $a, b \in \operatorname{supp} Q$, $\mathrm{v}(a^*) \neq \mathrm{v}(b)$, $a \geq b$ and $b^* \geq a^*$.

The *Q–inversion* with respect to $a \in \operatorname{supp} Q$ is $N_a N_{a^*}$. Note that a standard Q–transvection is not necessarily a transvection, but it is in some sense the closest thing to a transvection that preserves Q. In the case that Γ is a complete graph, the standard dominated Q–transvections are simply the standard symplectic transvections from classical linear algebra.

Claim 4.3   *The Q–transvections and Q–inversions preserve Q.*

Proof   Note the following computations:

$$N_a N_{a^*} \cdot (a \wedge a^*) = (-a) \wedge (-a^*) = a \wedge a^*$$
$$E_{a,a^*} \cdot (a \wedge a^*) = a \wedge (a^* + a) = a \wedge a^*$$
$$E_{a,b} E_{b^*,a^*}^{-1} \cdot (a \wedge a^* + b \wedge b^*) = E_{a,b} \cdot (a \wedge a^* + b \wedge b^* - a \wedge b^*)$$
$$= a \wedge a^* + b \wedge b^*$$

The claim follows immediately.   □

Let $G < \operatorname{Aut} H_\Gamma$ be the group:

$$G = \langle \{ E_{a,b} \mid a \in X, \ b \in \operatorname{supp} Q \text{ and } a \geq b \} \cup \{ N_b \mid b \in \operatorname{supp} Q \} \rangle$$

This is the image under the homology representation of the subgroup of $\operatorname{Aut}^0 A_\Gamma$ that fixes each element of $(\operatorname{supp} w_0)^{\pm 1}$.

This subsection is devoted to the proof of the following theorem.

Theorem 4.4   *The subgroup of G stabilizing Q is generated by the standard dominated Q–transvections and Q–inversions.*

The basis X lets us think of End H_Γ as matrices; in particular, it allows us to identify End H_Γ with $\otimes^2 H_\Gamma$ and gives us a transpose operation. Express End H_Γ as three-by-three block matrices, with respect to the decomposition of H_Γ as

$$\langle \operatorname{supp} w_0 \rangle \oplus \langle a_{k+1}, \dots, a_g \rangle \oplus \langle a_{k+1}^*, \dots, a_g^* \rangle.$$

Define $J \in \operatorname{End} H_\Gamma$ by:

$$J := \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -I_{g-k} \\ 0 & I_{g-k} & 0 \end{pmatrix}$$

Then J is the image of Q under the map:

$$\Lambda^2 H_\Gamma \hookrightarrow \otimes^2 H_\Gamma \xrightarrow{\cong} \operatorname{End} H_\Gamma$$

Note that for any $A \in \operatorname{Aut} H_\Gamma$, we have $A \cdot Q = Q$ if and only if $AJA^\top = J$. Also note that for any $a \in \operatorname{supp} Q$, we have $Ja = a^*$. Let $H_Q < H_\Gamma$ be generated by the image of supp Q.

Lemma 4.5   *If $A_0 \in G$ and A_0 fixes Q, then A_0 leaves H_Q invariant.*

Proof   From the definition of G we have that for some matrices A, B, C, D, E, F:

$$A_0 = \begin{pmatrix} I_{2k} & E & F \\ 0 & A & B \\ 0 & C & D \end{pmatrix}$$

Since $A_0 J A_0^\top = J$, we can deduce that:

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} \cdot \begin{pmatrix} -F^\top \\ E^\top \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

But since the matrix $\begin{pmatrix} A & B \\ C & D \end{pmatrix}$ is a diagonal block of a block-upper-triangular matrix, it is invertible, and therefore $F = 0$ and $E = 0$. □

By virtue of Lemma 4.5, we restrict our entire argument from H_Γ to H_Q. We also use the symbols Q and J to represent their respective restrictions to H_Q. Note that J restricted to H_Q is invertible. The element Q is a standard symplectic form, when considered as an element of $\Lambda^2 H_Q$.

Lemma 4.6   *If $\alpha \in \mathrm{Aut}^0 A_\Gamma$, then for any $a, b \in X$, we have $\langle \alpha_* b, a \rangle \neq 0$ implies $a = b$ or $a \geq b$.*

Proof   We induct on the length of α in terms of the generators of $\mathrm{Aut}^0 A_\Gamma$. The assertion is obvious if the length is zero. Assume it is true for β and that $\alpha = \beta \gamma$ where γ is one of the generators of $\mathrm{Aut}^0 A_\Gamma$. If γ is an inversion or a partial conjugation, then lemma follows. Suppose $\gamma = \tau_{c,d}$ for some $c, d \in L$ with $c \geq d$. Then $\langle \alpha_* b, a \rangle \neq 0$ implies either that $\langle \beta_* b, a \rangle \neq 0$ or that $\mathrm{v}(a) = \mathrm{v}(c)$ and $\langle \beta_* b, d \rangle \neq 0$. In the first case, the lemma follows. In the second case, we have $a \geq d$ and also $d \geq b$ by inductive assumption. $\square$

We define a binary relation on $\mathrm{supp}\, Q$ called Q–*domination*: $a \overset{\cdot}{\geq} b$ if $\mathrm{v}(a) \neq \mathrm{v}(b^*)$ and $a \geq b$ and $b^* \geq a^*$, or if $\mathrm{v}(a) = \mathrm{v}(b^*)$ and $a \geq b$. We define Q–*domination equivalence* by: $a \overset{\cdot}{\sim} b$ if $a \overset{\cdot}{\geq} b$ and $b \overset{\cdot}{\geq} a$. Note that we have a Q–transvection sending b to $b + a$ only if $a \overset{\cdot}{\geq} b$.

Lemma 4.7   *If $A \in G$ and A fixes Q, then for any $a, b \in \mathrm{supp}\, Q$, we have that $\langle Ab, a \rangle \neq 0$ implies $a = b$ or $a \overset{\cdot}{\geq} b$.*

Proof   If $\langle Ab, a \rangle \neq 0$, then by Lemma 4.6, we have $a \geq b$. Since $AJA^\top = J$, we know $A = J(A^\top)^{-1} J^{-1} = J^\top (A^{-1})^\top J$. So

$$0 \neq \langle J^\top (A^{-1})^\top J b, a \rangle = \langle b^*, A^{-1} a^* \rangle$$

which implies (again by Lemma 4.6) that $b^* \geq a^*$. $\square$

Now we will reassign the indices for our basis for H_Q. Assume we have labeled some vertices $\{x_1, \ldots, x_i, y_1, \ldots, y_i\} \subset (\mathrm{supp}\, Q)^{\pm 1}$. Then we choose x_{i+1} to be Q–domination maximal among the elements of $\mathrm{supp}\, Q$ not yet labeled as x_j or y_j. Set $y_{i+1} = J x_{i+1}$. By construction, we deduce that $\{\mathrm{v}(x_1), \ldots, \mathrm{v}(x_i), \mathrm{v}(y_1), \ldots, \mathrm{v}(y_i)\}$ does not contain $\mathrm{v}(y_{i+1})$. We proceed this way until we have constructed a basis.

We will now prove Theorem 4.4 by exhibiting a row reduction algorithm. This algorithm will differ from the usual integral symplectic row reduction algorithm in that we have to check at each step that the Q–domination relation allows us to use a given Q–transvection.

Proof of Theorem 4.4   Let $A \in G$ fix Q. By Lemma 4.5, we think of A as being in $\mathrm{Aut}\, H_Q$. Assume inductively that we have already row-reduced A by applying

standard Q–transvections and inversions to get a matrix A_i (for $1 \le i \le g - k$) such that for $j < i$, we have

$$(4\text{-}2) \qquad\qquad\qquad\qquad A_i x_j = x_j.$$

Since A is symplectic, A_i is symplectic, so $A_i = J^\top (A_i^{-1})^\top J$, and

$$(4\text{-}3) \qquad\qquad \langle A_i x_i, y_j \rangle = -\langle y_i, A_i^{-1} x_j \rangle = -\langle y_i, x_j \rangle = 0$$

for any $j < i$. If we have any $j \ge i$ with $\langle A_i x_i, x_j \rangle \ne 0$, then by Lemma 4.7, we have $x_j \gtrdot x_i$; since x_i is maximal among $\{x_i, \dots, x_{g-k}\}$, we know $x_i \sim x_j$. Similarly, if $j \ge i$ with $\langle A_i x_i, y_j \rangle \ne 0$, we know $x_i \mathbin{\dot\sim} y_j$.

Step 1 Consider all the indices $j \ge i$ such that both $\langle A_i x_i, x_j \rangle \ne 0$ and $\langle A_i x_i, y_j \rangle \ne 0$. For any such j, we have $x_j \mathbin{\dot\sim} y_j$, and by repeatedly applying the row operations E_{x_j, y_j}, E_{y_j, x_j} and their inverses (which are Q–transvections) according to the Euclidean algorithm, we can reduce A_i to a matrix $\widehat{A}_i$ in which either $\langle \widehat{A}_i x_i, x_j \rangle = 0$ or $\langle \widehat{A}_i x_i, y_j \rangle = 0$. By repeating this step for all such indices j, we assume we have reduced A to A_i' in which for each $j \ge i$, either $\langle A_i' x_i, x_j \rangle = 0$ or $\langle A_i' x_i, y_j \rangle = 0$. Note that these operations do not affect the columns of A_i before column of x_i, so Equation (4-2) and Equation (4-3) both still hold with A_i' instead of A_i.

Step 2 We find an element $a \in \{x_i, \dots, x_{g-k}, y_i, \dots, y_{g-k}\}$ maximizing $|\langle A_i' x_i, a \rangle|$ for a in this set. Since the determinant of A_i' is nonzero, we can deduce from the form of A_i' that this maximum is nonzero. If this a is the unique such element making this value nonzero, we move on to the next step. Otherwise, there is some other $b \in \{x_i, \dots, x_{g-k}, y_i, \dots, y_{g-k}\}$ with $|\langle A_i' x_i, b \rangle| \ne 0$. Since these matrix entries are nonzero, we know that $a \mathbin{\dot\sim} b$. Since by the first step, we know that $\langle A_i' x_i, a^* \rangle = \langle A_i' x_i, b^* \rangle = 0$, we know that the row operation $E_{b,a} E_{b^*,a^*}^{-1}$ and its inverse only change the column of x_i in A_i' by adding plus or minus the b–entry to the a–entry. Further, this does not alter the column of x_j in A_i' for any $j < i$. This step reduces either the maximum of $|\langle A_i' x_i, a \rangle|$ for $a \in \{x_i, \dots, x_{g-k}, y_i, \dots, y_{g-k}\}$, or it reduces the number of elements realizing this maximum absolute value. Either way, by repeatedly applying this step, we arrive at a matrix A_i'' such that there is a unique $a \in \{x_i, \dots, x_{g-k}, y_i, \dots, y_{g-k}\}$ with $|\langle A_i'' x_i, b \rangle| \ne 0$. Again, the Equation (4-2) and Equation (4-3) both still hold with A_i'' instead of A_i.

Step 3 We now have a unique $a \in \{x_i, \dots, x_{g-k}, y_i, \dots, y_{g-k}\}$ with $\langle A_i'' x_i, a \rangle \ne 0$. By the form of A_i'' and the fact that its determinant is 1, we deduce that $|\langle A_i'' x_i, a \rangle| = 1$.

If $v(a) \ne v(x_i)$, then we know $a \sim x_i$. In this case, if $v(a) \ne v(y_i)$ we can apply the product of Q–transvections

$$(E_{a, x_i} E_{x_i^*, a^*}^{-1})(E_{x_i, a}^{-1} E_{a^*, x_i^*})(E_{a, x_i} E_{x_i^*, a^*}^{-1})$$

which sends x_i to a, a to $-x_i$, y_i to a^* and a^* to $-y_i$ while fixing all other elements of our basis. Otherwise, $\mathrm{v}(a) = \mathrm{v}(y_i)$ and we can apply the product of Q–transvections

$$E_{y_i,x_i} E_{x_i,y_i}^{-1} E_{y_i,x_i}$$

which sends x_i to y_i and y_i to x_i^{-1} while fixing all other elements of our basis. In any event, perhaps by applying some Q–transvections, we may assume that $\mathrm{v}(a) = \mathrm{v}(x_i)$. Possibly after applying a Q–inversion, we may assume that $\langle A_i'' x_i, x_i \rangle = 1$. So for $j \geq i$, we have $\langle A_i'' x_i, y_j \rangle = 0$ and $\langle A_i'' x_i, x_j \rangle$ is 0 if $i \neq j$ and 1 if $i = j$.

Step 4 For each $j < i$ with $\langle A_i'' x_i, x_j \rangle \neq 0$, we know that $x_j \overset{\cdot}{\geq} x_i$, and we may apply the row operation $E_{x_j,x_i} E_{y_i,y_j}^{-1}$. Since all of the other relevant entries in the column of x_i are zero, the only effect of this operation is to add 1 to the x_j–entry. Of course, by applying some power of this operation, we can delete this entry. By applying this step repeatedly, we arrive at a matrix that satisfies the recursion hypotheses for $i + 1$ and we can go back to step 1.

We recurse through these steps $g - k$ times and arrive at a matrix A_{g-k+1} satisfying Equation (4-2) and Equation (4-3) for $i = g - k + 1$. However, these conditions imply that A_{g-k+1} is the identity matrix.

Since we reduced the arbitrary A to the identity matrix by repeatedly applying Q–transvections and Q–inversions, we have shown that these elements generate the stabilizer of Q in G. $\qquad\square$

4.3 Automorphisms fixing a surface relator

Recall the bijection $*\colon L \to L$ with $(a^*)^* = a^{-1}$ for all $a \in L$ and the surface relator

$$w_0 = [a_1, b_1] \cdots [a_k, b_k].$$

Note that $|w_0| = 4k$. In this subsection, we assume that $|w_0| > 0$. Note that if $a \in \operatorname{supp} w_0$, then either $(a^*)^{-1} a a^*$ or $a^* a^{-1} (a^*)^{-1}$ is a subsegment of w_0.

From here on, we will use symbols like w to refer to a word or the group element it determines, and we will use $[w]$ to refer to the cyclic word determined by w or to the conjugacy class of w.

The goal of this subsection is to prove that we can peak-reduce an arbitrary automorphism in $\operatorname{Aut} A_\Gamma$ (not just in $\langle \Omega_\ell \rangle$) if we are only reducing peaks with respect to $[w_0]$ (Theorem 4.17). In order to do this, we split an automorphism fixing $[w_0]$ into its long-range and short-range parts, and we will analyze this short-range part (Sublemma 4.14 and Lemma 4.15). Once we understand the short-range part, we will

be able to absorb all of the peak-forming short-range elements into general Whitehead automorphisms (Lemma 4.16).

We define the *pure long-range Whitehead automorphisms* Ω_ℓ^0 to be $\Omega_\ell \cap \text{Aut}^0 A_\Gamma$. We will use a slight refinement of part (3) of Theorem 2.1: if $[w]$ is a conjugacy class and $\alpha \in \langle \Omega_\ell^0 \rangle$, then α can be peak-reduced with respect to $[w]$ by elements of Ω_ℓ^0 (see Day [6, Remark 3.22]).

Our first goal is to show that the action of $\text{Aut}^0 A_\Gamma$ cannot shorten $[w_0]$ or shrink its support. This relies strongly on the structure of w_0, which we exploit through the following two sublemmas.

Sublemma 4.8 *Suppose C is a nonempty adjacent domination equivalence class of Γ and suppose γ is a product of transvections and partial conjugations with multipliers in C. Let $j = |C \cap \text{supp } w_0|$ and let $m = (k - j)$. There are letters $f_1, \ldots, f_j \in (\text{supp } w_0 \cap C)^{\pm 1}$, $x_1, \ldots, x_j$, $c_1, \ldots, c_m$ and $d_1, \ldots, d_m$ in $L - C^{\pm 1}$ and words $u_i = \gamma(f_i)$, $r_1, \ldots, r_m$, $s_1, \ldots, s_m$ and $t_1, \ldots, t_m$ in $C^{\pm 1}$ such that $\gamma([w_0])$ is represented by a cyclic word given as a graphically reduced product of the words*

$$x_1 u_1 x_1^{-1}, \ldots, x_j u_j x_j^{-1}$$

with the words

$$c_1 r_1 d_1 s_1 c_1^{-1} t_1 d_1^{-1}, \ldots, c_m r_m d_m s_m c_m^{-1} t_m d_m^{-1}$$

and some elements of $C^{\pm 1}$ in some order. Further, if $j > 0$ then these $\{u_i\}_i$ are all nontrivial, and their product is nontrivial.

Proof We will prove this statement by induction on the length of γ as a product of Whitehead automorphisms. First we discuss the base case. For a factor $[a_i, b_i]$ of w_0, not both of a_i and b_i are in C since a_i and b_i do not commute. If a given a_i is in $C^{\pm 1}$, we set the next available $x_p = b_i$ and set $u_p = f_p = a_i^{-1}$. Similarly, if b_i is in $C^{\pm 1}$, we set the next available $x_p = a_i$ and set $u_p = f_p = b_i$. For each i with both $\text{v}(a_i), \text{v}(b_i) \notin C$, we set the next available $c_p = a_i$ and set $d_p = b_i$. We take each r_p, s_p and t_p to be the empty word. This proves the base case $|\gamma| = 0$.

Now suppose that w satisfies the conclusions of the lemma and let α be a Whitehead automorphism with multiplier $a \in C^{\pm 1}$, so that $\gamma = \alpha \gamma'$ for some γ'. Then for each element x of $X - C$, $\alpha(x)$ contains a single instance of x, and $\text{supp } \alpha(x) \subset \{\text{v}(a), x\}$. Then the same choices of $\{x_i\}_i$, $\{c_i\}_i$ and $\{d_i\}_i$ elements will work. For each i, it is possible to choose new words r_i, s_i and t_i that will work based on the old words and $\alpha(c_i)$ and $\alpha(d_i)$. Note that regardless of what $\alpha(x_i)$ is, $\alpha(x_i u_i x_i^{-1})$ is $x_i \alpha(u_i) x_i^{-1}$ or $a^{-1} x_i \alpha(u_i) x_i^{-1} a$. This means that our $\alpha(u_i)$ will work as our new u_i, and since

$u_i = \gamma'(f_i)$, we have $\alpha(u_i) = \gamma(f_i)$. This means that we can write $\alpha(w)$ in the desired form.

Note that since each original f_i is nontrivial, each $\gamma(f_i)$ is nontrivial, and since the product $f_1 \cdots f_j$ is nontrivial, the product $\gamma(f_1) \cdots \gamma(f_j)$ is nontrivial. $\qquad\square$

Sublemma 4.9 *Suppose C is an adjacent domination equivalence class of Γ with $|C| > 1$, and suppose γ is a product of transvections and partial conjugations with multipliers in C. Then if $\alpha \in \Omega_\ell$ shortens $\gamma([w_0])$, then $\alpha = (A, a)$ for some $a \in C^{\pm 1}$.*

Further, no $\alpha \in \Omega_\ell$ can shorten $[w_0]$.

Proof We will prove both statements at once by supposing that either $C = \varnothing$ (and γ is trivial) or $|C| > 1$. If $\alpha \in \Omega_\ell$ shortens $\gamma([w_0])$, then α is not a permutation automorphism, so suppose $\alpha = (A, a)$. Suppose for contradiction that $v(a) \notin C$. Let w be a representative of $\gamma([w_0])$ of the form given in Sublemma 4.8 (or $w = w_0$ if $C = \varnothing$). The conjugacy class of w maps to the trivial element of $\operatorname{Aut} H_\Gamma$. This means that every element of $\operatorname{supp} w$ appears an even number of times in w, half with positive exponent and half with negative exponent. Since $v(a) \notin C$, we know by Sublemma 4.8 that a appears only twice in w. So we have v_1, v_2 words in $L - \{a, a^{-1}\}$ such that $w = v_1 a v_2 a^{-1}$ as graphically reduced cyclic words. Then since α shortens w, it must delete both the instance of a and the instance of a^{-1} in w without introducing any new instances of $a^{\pm 1}$. Then $\alpha(v_1) = v_1$ and $\alpha(v_2) = a^{-1} v_2 a$ (since α multiplies generators by a only on the right).

In the case $|C| > 1$, we have distinct $b, c \in C^{\pm 1}$. Suppose $a \geq b$. Then $a \in \operatorname{lk}_L(c)$, and therefore $a \in \operatorname{lk}_L(b)$. So either a commutes with every element of C or a does not dominate any element of C.

With notation as in Sublemma 4.8, we first suppose that $v(a)$ is a $v(c_i)$ or a $v(d_i)$. This is the only possibility if $C = \varnothing$. We suppose that $v(a) = v(c_i)$, since the case that $v(a) = v(d_i)$ is parallel. Then possibly by swapping (A, a) with $(L - A - \operatorname{lk}_L(a), a^{-1})$, we assume that $a = c_i$. Then $v_2 = r_i d_i s_i$. If $C = \varnothing$, of course, our words r_i, s_i and t_i are all empty. If $r_i \neq 1$ or $s_i \neq 1$ and a does not commute with the elements of C, then to delete the instances of $a^{\pm 1}$ in $a r_i$ and $s_i a^{-1}$, α must conjugate C. Then in any event, either to delete existing instances or to avoid introducing new instances of $a^{\pm 1}$, d_i and d_i^{-1} must be in A (note that $d_i \notin \operatorname{lk}_L(a)$ since $a = c_i$). We know that $t_i d_i^{-1}$ is an initial segment of v_1. If $t_i = 1$, then we already have a contradiction, since d_i is then the first letter of v_1 and we have $\alpha(v_1) \neq v_1$. If $t_i \neq 1$, a does not commute with t_i, and to avoid introducing an extra instance of a between t_i and d_i^{-1}, α must conjugate C. But then $\alpha(t_i) = a^{-1} t_i a^{-1}$, and we cannot have $\alpha(v_1) = v_1$, which is

a contradiction. So $v(a)$ is not a $v(c_i)$ or a $v(d_i)$. Note that in the case $C = \varnothing$, the proof is finished.

Then it must be that $v(a) = v(x_i)$ for some i. For the rest of the proof, we assume $|C| > 1$. We suppose without loss of generality that $a = x_i$. Then $v_2 = u_i$, a word in $C^{\pm 1}$. If a commutes with the elements of C, then since α is long-range, α fixes u_i, and we have $\alpha(v_2) = v_2$, which is a contradiction. So suppose a does not dominate any element of C. Then α sends every element of C to its conjugate by a. Since u_i is nontrivial and w maps to the trivial element of H_Γ, we know that there are some elements of C in v_1. So there is a subsegment v_3 of v_1 such that $\alpha(v_3) = a^{-1} v_3 a$. Let v_4 be the longest subsegment of v_1, containing v_3, such that $\alpha(v_4) = a^{-1} v_4 a$. Since $\alpha(v_1) = v_1$, we know that if we delete v_4 from v_1 we get two subsegments.

By Sublemma 4.8, the letter furthest to the left in this right subsegment of v_1 must be an element of $C^{\pm 1}$, or an x_i, or a c_i. If it is an element of $C^{\pm 1}$, then it is conjugated by α, contradicting the definition of v_4. If it is an x_i, then this x_i is in a subsegment $x_i u_i x_i^{-1}$. Since v_4 maps to $a^{-1} v_4 a$, we know that x_i^{-1} must be in A or else that $x_i \in \mathrm{lk}_L(a)$. If $x_i \notin \mathrm{lk}_L(a)$, then the x_i^{-1} on the right adds an instance of a, contradicting the definition of v_4. If $x_i \in \mathrm{lk}_L(a)$, then since the u_i is nontrivial, it is conjugated by a, also contradiction the definition of v_4. If this letter is a c_i, there are several cases. If $c_i, d_i \in \mathrm{lk}_L(a)$, then either the a from $\alpha(v_4)$ commutes past our $c_i r_i d_i s_i c_i^{-1} t_i d_i^{-1}$, or one of the r_i, s_i or t_i is nontrivial and an a is introduced by conjugation. If d_i is in $\mathrm{lk}_L(a)$ but c_i is not, then to avoid introducing an a, we have $c_i^{-1} \in A$ and another a is introduced either by the c_i^{-1} or the t_i. If c_i is in $\mathrm{lk}_L(a)$ but d_i is not, then either the a from v_4 or from r_i must be cancelled by an a^{-1} from d_i, so $d_i^{-1} \in A$ and the final d_i^{-1} introduces an extra a. If both $c_i, d_i \notin \mathrm{lk}_L(a)$, then to cancel the a from v_4, $c_i^{-1} \in A$; to cancel the a from c_i^{-1} or t_i, we have $d_i \in A$; to cancel the a from d_i or from s_i, we have $c_i \in A$; and to cancel the a from c_i or r_i, we have $d_i^{-1} \in A$. This means that d_i^{-1} introduces an extra a at the end. In any event we contradict the definition of v_4 if the letter in v_1 right after v_4 is a c_i. So v_4 must extend to the right edge of v_1, a contradiction. □

Lemma 4.10 *Suppose C is the domination equivalence class of an element $c \in X$. Then if $\gamma \in \mathrm{Aut}^0 A_\Gamma$, we have $|C \cap \mathrm{supp}\, \gamma([w_0])| \geq |C \cap \mathrm{supp}\, w_0|$.*

Proof By Theorem 2.1, we can write $\gamma = \alpha\beta$ where $\alpha \in \langle \Omega_\ell^0 \rangle$ and $\beta \in \langle \Omega_s \rangle$. By Theorem 2.1, we can write $\beta = \beta' \beta_C$, where β_C is a product of short-range transvections with multipliers in C and β' is a product of short-range transvections with multipliers not in C. Again by Theorem 2.1, we can rewrite $\alpha\beta'$ as $\beta''\alpha'$ for some $\alpha' \in \langle \Omega_\ell^0 \rangle$ and some $\beta'' \in \langle \Omega_s \rangle$. Further, by the form of the sorting substitutions in

Definition 3.2 of Day [6], we know that β'' is also a product of short-range transvections with multipliers not in C.

We have $\gamma = \beta''\alpha'\beta_C$. The automorphism β'' cannot remove any instances of elements of C from the support of a word because it can be written as a product of transvections whose multipliers are not in C. We will prove the lemma by showing that β_C cannot send w_0 to a word containing fewer elements of C, and then by showing that α' cannot remove elements of C from $\beta_C([w_0])$.

Consider $(\beta_C)_* \in \operatorname{Aut} H_\Gamma$ as a matrix with respect to the generators of H_Γ given by the image of X. By Day [6, Corollary 3.11], we know that the block of this matrix sending the image of C to itself is invertible. Then:

$$\left| \bigcup \{ C \cap \operatorname{supp} \beta_C(a) \mid a \in C \cap \operatorname{supp} w_0 \} \right| \geq |C \cap \operatorname{supp} w_0|$$

Suppose temporarily that C is an adjacent domination equivalence class. If $\operatorname{v}(a) \in C \cap \operatorname{supp} w_0$, we know that $\operatorname{v}(a^*) \notin C$ since $a^* \notin \operatorname{lk}_L(a)$. Then $\operatorname{v}(a^*) \notin \operatorname{supp} \beta_C(x)$ for any $x \neq (a^*)^{\pm 1}$, and since $a^* \notin \operatorname{lk}_L(a)$, we also know $\beta_C(a^*) = a^*$. If C is a nonadjacent domination equivalence class, then $\beta_C = 1$, and it follows in both cases that $\beta_C(a^*) = a^*$ for every a with $\operatorname{v}(a) \in C \cap \operatorname{supp} w_0$.

Consider the representative w for $\beta_C([w_0])$ from Sublemma 4.8. For each element $a \in C \cap \operatorname{supp} w_0$, $\operatorname{v}(a^*)$ is one of the $\operatorname{v}(x_i)$ elements and $\beta_C(a)$ is the corresponding u_i. In particular, there are at least as many elements of C appearing in subsegments of w of the form $(a^*)^{-1}\beta_C(a)a^*$ or $a^*\beta_C(a)^{-1}(a^*)^{-1}$, for $a \in C^{\pm 1}$ and $a^* \notin C^{\pm 1}$ as there are elements of $C \cap \operatorname{supp} w_0$.

By Theorem 2.1, α' has a factorization by elements of Ω_ℓ^0 that is peak-reduced with respect to $\beta_C([w_0])$. This factorization may include permutations, but these permutations preserve C (because they are in $\operatorname{Aut}^0 A_\Gamma$), so the only way to remove any extra instances of elements of C from w is to decrease its length. Then peak reduction implies that the elements of this factorization shorten w immediately and keep shortening it until all the excess instances of elements of C have been removed. If C is an adjacent domination equivalence class and $|C| > 1$, then by Sublemma 4.9, each one of these shortening automorphisms has multiplier in C, and again by Sublemma 4.8, we see that these shortening automorphisms do not remove any elements from $C \cap \operatorname{supp} w_0$. If $|C| = 1$ or C is a nonadjacent domination equivalence class, then $\beta_C = 1$ and $\beta_C([w_0]) = [w_0]$. Then by Sublemma 4.9, no element of Ω_ℓ can shorten $\beta_C([w_0])$, and therefore α' cannot remove any elements from $C \cap \operatorname{supp} w_0$. $\square$

Corollary 4.11 *For each $\gamma \in \operatorname{Aut}^0 A_\Gamma$, we have $|\gamma([w_0])| \geq |w_0|$.*

Proof We know from Lemma 4.10 that $|\operatorname{supp}\gamma([w_0])| \geq 2k$. Since $\gamma([w_0])$ maps to the trivial element of H_Γ, each element appears at least twice (once with positive and once with negative exponent). So $|\gamma([w_0])|$ is at least $4k$, the length of w_0. $\square$

Now we will analyze the short-range part of an automorphism fixing $[w_0]$.

Definition 4.12 A short-range transvection $\tau_{a,b}$ is w_0–*irrelevant* if $a \in \operatorname{lk}_L(b^*)$. An automorphism is w_0–*irrelevant* if it is a product of w_0–irrelevant transvections.

Remark 4.13 Suppose $e(a,b^*)$ and $a \geq b$. Then

$$\tau_{a,b}(bb^*b^{-1}(b^*)^{-1}) = bab^*a^{-1}b^{-1}(b^*)^{-1} = bb^*b^{-1}(b^*)^{-1}.$$

So $\tau_{a,b}$ fixes w_0. We call these automorphisms w_0–irrelevant because they are an obvious class of automorphisms fixing w_0.

Sublemma 4.14 *Let C be an adjacent domination equivalence class in X. Suppose $\alpha \in \langle \Omega_\ell^0 \rangle$, σ is a permutation automorphism that fixes C, γ is a product of short-range transvections with multipliers not in C, and β is a product of short-range transvections with multipliers in C, such that*

$$\gamma\sigma\alpha\beta([w_0]) = [w_0].$$

Then we can write β as $\beta'\iota$, where ι is w_0–irrelevant and β' is a product of short-range transvections with multipliers in C, none of which are w_0–irrelevant.

Proof First we note that the group

$$\big\langle \{\tau_{a,b} \mid v(b) \in \operatorname{supp} w_0 \text{ and } a \in C \cap \operatorname{lk}_L(b) \cap \operatorname{lk}_L(b^*)\} \big\rangle$$

is a normal subgroup of the group generated by short-range transvections with multipliers in C. If $b \in \operatorname{supp} w_0$ and $a \in C \cap \operatorname{lk}_L(b) \cap \operatorname{lk}_L(b^*)$, then $b \notin C$. If b were in C, then $b \sim a$ and $b \in \operatorname{lk}_L(b^*)$, a contradiction. So for $\tau_{c,d}$ a short-range transvection with $c \in C$, we have $v(c) \neq v(b)$, and either $\tau_{c,d}$ commutes with $\tau_{a,b}$, or $v(d) = v(a)$ and we apply $[\tau_{c,a},\tau_{a,b}] = \tau_{c,b}$. Since $c \geq a$, we have $c \in \operatorname{lk}_L(b)$, and $\tau_{c,b}$ is a member of the subgroup and the subgroup is normal.

So we can move w_0–irrelevant transvections $\tau_{a,b}$ with $v(b) \in \operatorname{supp} w_0$ to the right of any other transvections in a factorization of β, and therefore without loss of generality we may assume that β has a factorization in which the only w_0–irrelevant transvections that appear are the ones of the form $\tau_{a,b}$ where $v(b) \notin \operatorname{supp} w_0$.

Apply β to w_0 letter-by-letter and graphically reduce to get a cyclic word w. Then w is a representative of $\beta([w_0])$. Suppose there is some $b \in \operatorname{supp} w_0$ and $y \in C - \operatorname{supp} w_0$

with $y \in \operatorname{supp} \beta(b)$. By the form of β, we know that no element in $\operatorname{supp} \beta(b)$ commutes with b^*. Since b^* does not commute with b, we know b^* is not in C, and therefore each instance of b^* survives in w. No instance of y can be cancelled out of the subsegment $b^*\beta(b)^{-1}(b^*)^{-1}$ or $(b^*)^{-1}\beta(b)b^*$ of w (this is the image of the subsegment $b^*b^{-1}(b^*)^{-1}$ or $(b^*)^{-1}bb^*$ of w_0).

Since $y \notin \operatorname{supp} w_0$, either α or γ or σ must remove it. We know that σ fixes C, so σ cannot remove it. Also γ cannot remove y because γ can be written as a product of transvections whose multipliers are not y. We can peak-reduce α with respect to w. This peak-reduced factorization may have permutation automorphisms in it, but these will fix adjacent domination equivalence classes. So, there must be a sequence of long-range automorphisms, each of which progressively shortens w, which remove all instances of y. This is impossible: the b^* and $(b^*)^{-1}$ cannot be removed since removing one of them would change the class of the word in H_Γ and removing both would contradict Lemma 4.10; they cannot be moved without being removed since this would not shorten the word; and without moving or removing the b^* and $(b^*)^{-1}$ it is impossible to remove the instance of y between them. This is a contradiction, so we may assume that for $b \in \operatorname{supp} w_0$, we have $\operatorname{supp} \beta(b) \subset \operatorname{supp} w_0$.

This fact, together with Theorem 2.1, lets us deduce that β has a factorization by short-range transvections with multipliers in $C \cap \operatorname{supp} w_0$. Note that the subgroup

$$\left\langle \{\tau_{x,y} \mid x \in C \cap \operatorname{supp} w_0 \text{ and } y \notin \operatorname{supp} w_0\} \right\rangle$$

is normal in the group of short-range automorphisms with multipliers in $C \cap \operatorname{supp} w_0$. This is because for any $\tau_{x,y}$ with $x \in C \cap \operatorname{supp} w_0$ and $y \notin \operatorname{supp} w_0$, and any $\tau_{a,b}$ with $a \in C \cap \operatorname{supp} w_0$, either $\tau_{x,y}$ and $\tau_{a,b}$ commute or $\mathrm{v}(b) = \mathrm{v}(x)$ and we apply the identity $[\tau_{a,b}, \tau_{b,x}] = \tau_{a,x}$. Since this subgroup is normal, we can rewrite β with all the w_0–irrelevant transvections first. $\qquad\square$

The following lemma is the core reason that we are able to peak-reduce automorphisms fixing $[w_0]$, regardless of whether they are long-range or not.

Lemma 4.15 *Suppose $\alpha \in \langle \Omega_\ell^0 \rangle$ and $\beta \in \langle \Omega_s \rangle$ such that $\alpha\beta([w_0]) = [w_0]$. Then there is a permutation automorphism σ that leaves $\operatorname{supp} w_0$ invariant, a w_0–irrelevant automorphism $\iota \in \langle \Omega_s \rangle$, distinct elements $x_1, \ldots, x_r \in (\operatorname{supp} w_0)^{\pm 1}$ (with $x_i \neq x_j^{-1}$ for any i, j) and elements $y_1, \ldots, y_r \in (\operatorname{supp} w_0)^{\pm 1}$ with $x_i \geq y_i$ and $y_i \in \operatorname{lk}_L(x_i)$ such that*

$$\beta = \sigma \tau_{x_1, y_1} \cdots \tau_{x_r, y_r} \iota$$

and such that each x_i is domination-minimal among $\{x_i, x_{i+1}, \ldots, x_r\}$.

Proof Suppose that $C_1 \cup \cdots \cup C_m = X$ is the decomposition of X into adjacent domination equivalence classes. We assume that these sets are indexed such that if $a \in C_i$ and $b \in C_j$ with $a \in \mathrm{lk}_L(b)$, $a \geq b$ and $a \not\sim b$, then $i > j$. This assumption makes C_1 minimal and makes C_m maximal.

Inductively assume we have expressed β as

$$\beta' \sigma_{p-1} \delta \iota_{p-1}$$

where β' is a product of short-range transvections whose multipliers are in $C_p \cup \ldots \cup C_m$; the automorphism δ can be written as a product of short-range transvections with distinct multipliers in $C_1 \cup \ldots \cup C_{p-1}$, in domination-nondecreasing order; the automorphism ι_{p-1} is w_0–irrelevant; and σ_{p-1} is a permutation automorphism that is trivial outside of $C_1 \cup \ldots \cup C_{p-1}$. We will show that we can then do the same for p instead of $p-1$.

By Day [6, Corollary 3.11], we can rewrite β' as $\beta'' \beta_p$, where β_p is a product of short-range transvections whose multipliers are in C_p and β'' is a product of short-range transvections whose multipliers are in $C_{p+1} \cup \ldots \cup C_m$. Then β_p commutes with σ_{p-1}. We can also conjugate β_p across δ, as follows. Observe that if we have short-range transvections $\tau_{a,b}$ and $\tau_{c,d}$ with $c \in C_p$ and $a \in C_i$ for $i < p$, then the transvections commute unless $v(d) = v(a)$, in which case we have $[\tau_{c,a}, \tau_{a,b}] = \tau_{c,b}$. In any case, we do not change δ by conjugating these elements across it, and the new transvections we introduce have multipliers in C_p. As a result we can write β as

$$\beta'' \sigma_{p-1} \delta \beta'_p \iota_{p-1}$$

where β'_p is a product of transvections whose multipliers are in C_p.

Next we move β'' back across α and move δ across σ_{p-1} and α. Of course, this is possible by Theorem 2.1, but we also note that by Day [6, Equation (3-3)], if we introduce new short-range transvections through this process, they will have the same multipliers as those already in β'' and δ_{p-1}. So we can write

$$\alpha \beta'' \sigma_{p-1} \delta = \gamma \sigma_{p-1} \alpha'$$

where γ is a product of short-range transvections with multipliers not in C_p and $\alpha' \in \langle \Omega^0_\ell \rangle$.

Since ι_{p-1} is w_0–irrelevant, one can easily see that $\iota_{p-1}([w_0]) = [w_0]$. Then since $\alpha\beta = \gamma \sigma_{p-1} \alpha' \beta'_p \iota_p$, we have $\gamma \sigma_{p-1} \alpha' \beta'_p([w_0]) = [w_0]$. By Sublemma 4.14, we can write $\beta'_p \iota_{p-1}$ as $\beta''_p \iota_p$, where β''_p is a product of transvections with multiplier in C_p that are not w_0–irrelevant, and ι_p is a product of w_0–irrelevant transvections. In particular, we have $\beta = \beta'' \sigma_{p-1} \delta \beta''_p \iota_p$.

If we consider $(\beta_p'')_* \in \text{Aut}(H_\Gamma)$ as a matrix, we know that the block of $(\beta_p'')_*$ taking the image of C_p to itself is invertible. Further, since we have removed all the w_0–irrelevant automorphisms, we know that the block of $(\beta_p'')_*$ taking the image of $C_p \cap \text{supp } w_0$ to itself is invertible. Then there is a permutation σ_p' of $(C_p \cap \text{supp } w_0)^{\pm 1}$ such that $\sigma_p'(x)$ appears in $\beta_p''(x)$ to a positive power, for $x \in X$. We extend σ_p' by the identity outside of C_p to get a permutation of L; since C_p is an adjacent domination equivalence class this permutation extends to an automorphism of A_Γ. Let $\delta_p = (\sigma_p')^{-1}\beta_p''$, and let $\sigma_p = \sigma_{p-1}\sigma_p'$. Then each $x \in X$ appears in $\delta_p(x)$ to a positive power. Note that the hypotheses on δ imply that δ commutes with σ_p', and we have $\beta = \beta'' \sigma_p \delta \delta_p \iota_p$.

Let w be a word gotten by applying δ_p letter-by-letter to w_0 and graphically reducing. Since δ_p is free of w_0–irrelevant transvections in its factorization, for any $a \in \text{supp } w_0$, we know $\text{supp } \delta_p(a)$ does not contain any elements commuting with a^*. Then if we further suppose that $\text{supp } \delta_p(a) \neq \{a\}$, then a is adjacently dominated by an element of C_p, and we know that $a^* \notin C_p$ and we have $a^* \delta_p(a)^{-1}(a^*)^{-1}$ or $(a^*)^{-1}\delta_p(a)a^*$ as a subsegment in w. This is also true if $a \in C_p$. Of course, $\alpha\beta''\sigma_p\delta([w]) = [w_0]$. If we let $\alpha'' = (\sigma_p')^{-1}\alpha'\sigma_p' \in \langle \Omega_\ell \rangle$, then $\alpha\beta''\sigma_p\delta = \gamma\sigma_p\alpha''$. We know that γ cannot remove any instances of elements in C_p from a word, and σ_p can permute the elements of $C_p \cap \text{supp } w_0$ but cannot remove any. Therefore if the elements of $\text{supp } \delta_p(a)$ in w are removed by $\gamma\sigma_p\alpha''$, it must be α'' that removes them. We assume α'' to be peak-reduced with respect to $\delta_p([w_0])$, so there must be a sequence of long-range automorphisms that progressively shortens $\delta_p([w_0])$ and remove the extra instances of elements of C_p. However, we know we cannot alter the instance of $a^* \delta_p(a)^{-1}(a^*)^{-1}$ or $(a^*)^{-1}\delta_p(a)a^*$ in w by any such moves. If on the other hand $\delta_p(a) = a$, we know by Lemma 4.10 that $\sigma_p(a)$ survives to the final w_0.

Therefore for each element c appearing in $\delta_p(a)$ for any $a \in \text{supp } w_0$, the element $\sigma_p(c)$ appears in the final w_0 with at least the multiplicity with which c appears in $\delta_p(a)$. We know $\sigma_p(c) \in (\text{supp } w_0)^{\pm 1}$ if and only if $c \in (\text{supp } w_0)^{\pm 1}$. Therefore there cannot be any $a \in \text{supp } w_0$ with $\delta_p(a)$ containing any $x \in \text{supp } w_0$ to any power greater than 1 in absolute value, or with $\delta_p(a)$ containing any $x \notin \text{supp } w_0$ at all. Finally, if there are two distinct elements $c_1, c_2 \in \text{supp } w_0$ and some x with $x \in \text{supp } \delta_p(c_i)$ for $i = 1, 2$, then x appears in $c_i^* \delta_p(c_i)^{-1}(c_i^*)^{-1}$ or $(c_i^*)^{-1}\delta_p(c_i)c_i^*$ for $i = 1, 2$, and also in $x^* \delta_p(x)^{-1}(x^*)^{-1}$ or $(x^*)^{-1}\delta_p(x)x^*$ in w. So in this case, these three instances cannot be removed, and since two of them are both to a positive power or both to a negative power, there would be at least 4 instances of $\sigma_p(x)$ in w_0, which is impossible. So at most one element of $\text{supp } w_0$ maps to an element with a given x in its support under δ_p.

From this we deduce that the matrix $(\delta_p)_*$ has diagonal entries of 1, has off-diagonal entries of either ± 1 or 0, and has only trivial entries away from the rectangular block

sending elements dominated by C_p to the image of $C_p \cap \operatorname{supp} w_0$. Further, each row has at most one nonzero off-diagonal entry. Then the block sending the image of $C_p \cap \operatorname{supp} w_0$ to itself must be invertible; all together these conditions indicate that there is a reindexing of the basis that makes $(\delta_p)_*$ upper-triangular. An upper-triangular matrix where each row has at most one nonzero off-diagonal entry can be column reduced using each row operation at most once. By Theorem 2.1, we have factored δ_p as a product of short-range transvections with distinct multipliers in C_p.

Then one can easily see that $\beta = \beta'' \sigma_p (\delta \delta_p) \iota_p$ is a factorization satisfying the inductive hypothesis for the next step. The lemma follows. $\qquad\square$

Finally, we proceed to reducing peaks.

Lemma 4.16 *Let $\tau_{x,y} \in \Omega_s$ with $x, y \in \operatorname{supp} w_0$. Let $\alpha \in \langle \Omega_\ell^0 \rangle$. Suppose $\beta \in \langle \Omega_s \rangle$ is a product of transvections of the form $\tau_{a,b}$ for various $a \in \operatorname{supp} w_0$, $\mathrm{v}(a) \neq \mathrm{v}(x)$, such that x does not strictly dominate a. Further suppose that $\alpha \tau_{x,y} \beta([w_0])$ has the same length and support as w_0.*

Then we can find $(B, x) \in \Omega$ and $\alpha', \alpha'' \in \langle \Omega_\ell^0 \rangle$ such that $\alpha \tau_{x,y} \beta = \alpha'(B, x) \alpha'' \beta$ and $\alpha'' \beta([w_0])$ and $(B, x)\alpha'' \beta([w_0])$ have the same length and support as w_0.

Proof By Theorem 2.1, we peak-reduce α with respect to $\tau_{x,y} \beta([w_0])$. Then since $[w_0]$ is of minimal length in its $\operatorname{Aut} A_\Gamma$–orbit by Corollary 4.11, we have $\alpha = \gamma_1 \cdots \gamma_p \delta_1 \cdots \delta_q$, for some $\gamma_1, \ldots, \gamma_p, \delta_1, \ldots, \delta_q \in \Omega_\ell$, where each γ_i leaves the length of the word the same, and each δ_i shortens the word. More precisely, if $|\tau_{x,y} \beta([w_0])| = |w_0|$, then $q = 0$, if $q > 0$ then for each i we have

$$|\delta_i \cdots \delta_q \tau_{x,y} \beta([w_0])| < |\delta_{i+1} \cdots \delta_q \tau_{x,y} \beta([w_0])|$$

and if $p > 0$ then for each i we have

$$|\gamma_i \cdots \gamma_p \delta_1 \cdots \delta_q \tau_{x,y} \beta([w_0])| = |\gamma_{i+1} \cdots \gamma_p \delta_1 \cdots \delta_q \tau_{x,y} \beta([w_0])|.$$

Since β is a product of transvections with multipliers in $\operatorname{supp} w_0$, we know that $\operatorname{supp} \tau_{x,y} \beta([w_0]) \subset \operatorname{supp} w_0$. By Lemma 4.10, we know that they are equal. Since each δ_i decreases length, we know that

$$\operatorname{supp} \delta_i \cdots \delta_q \tau_{x,y} \beta([w_0]) \subset \operatorname{supp} \delta_{i+1} \cdots \delta_q \tau_{x,y} \beta([w_0]).$$

Again from Lemma 4.10, we know $\operatorname{supp} \delta_i \cdots \delta_q \tau_{x,y} \beta([w_0]) = \operatorname{supp} w_0$ for each i.

Temporarily fix i. The automorphism $\delta_i = (A, a)$ for some $a \in L$. Since δ_i decreases length, we know that $\mathrm{v}(a) \in \operatorname{supp} w_0$. By Lemma 4.10, we know δ_i cannot remove

all the instances of $a^{\pm 1}$ from $\delta_{i+1} \cdots \delta_q \tau_{x,y} \beta([w_0])$, so there must be extra instances of a in $\delta_{i+1} \cdots \delta_q \tau_{x,y} \beta([w_0])$. These extra instances must have been put there by $\tau_{x,y} \beta$ (since the other δ_j automorphisms are length-decreasing), so we deduce that either $v(a) = v(x)$, or that there is some $z \in \operatorname{supp} w_0$ such that $a \in \operatorname{supp} \beta(z)$. By the hypotheses on β, this tells us that if $v(a) \neq v(x)$, then x does not strictly dominate a.

Now we consider what happens when we try to move $\tau_{x,y}$ to the left across $\delta_i = (A, a)$. From Day [6, Lemma 3.4], we know that they commute (at least in Out A_Γ) unless $v(a) = v(y)$. Without loss of generality we temporarily assume $a = y$. In this case, conjugating $\tau_{x,y}$ across δ_i introduces a short range element $s((A - a + x, x))$ and a long-range element $\ell((A - a + x, x))$. However, since x does not strictly dominate a and $a = y$, we know $x \geq a$ and therefore $x \sim a$. If $x \sim a$, then the element $s((A - a + x, x)) = 1$. So in any case, we add at most a single new long-range element (working in Out A_Γ) and no new short-range elements. In returning to Aut A_Γ it is possible that we introduce an inner automorphism, which is a product of long-range automorphisms. So we have shown that there is an element $\phi_i \in \langle \Omega_\ell^0 \rangle$ such that $\delta_i \tau_{x,y} = \tau_{x,y} \phi_i$.

We rewrite $\delta_1 \cdots \delta_p \tau_{x,y}$ as $\tau_{x,y} \phi_1 \cdots \phi_p$. If

$$|\tau_{x,y} \phi_1 \cdots \phi_p \beta([w_0])| = |\phi_1 \cdots \phi_p \beta([w_0])|$$

then we are done; if we set $(B, x) = \tau_{x,y}$, set $\alpha' = \gamma_1 \cdots \gamma_q$ and set $\alpha'' = \phi_1 \cdots \phi_p$ then the conclusions hold. So assume $\tau_{x,y}$ decreases the length of the word. From the setup, we know that $\tau_{x,y} \phi_1 \cdots \phi_p \beta([w_0])$ has the same length and support as w_0. Then we know that y and y^{-1} both appear only once in $\phi_1 \cdots \phi_p \beta([w_0])$. This means that $\tau_{x,y}$ decreases the length by 2, removing an instance of x and x^{-1} each. By the form of β, we know that $\beta([w_0])$ only has a single x and a single x^{-1}, so $\phi_1 \cdots \phi_p$ must increase the number of instances of x. We have a word $\psi_1 \cdots \psi_r$ in Ω_ℓ^0 that is a peak-reduced factorization of $\phi_1 \cdots \phi_p$ with respect to $\beta([w_0])$. Some automorphism ψ_i adds an extra instance of x and in doing so increases the length by 2. Since the factorization is peak-reduced, this automorphism must be $\psi_1 = (B', x)$ (without loss of generality, we assume the multiplier is x and not x^{-1}, since $\tau_{x,y} = \tau_{x^{-1}, y^{-1}}$). We set $B = B' \cup \{y\}$ to get $\tau_{x,y}(B', x) = (B, x) \in \Omega$. Since $|\psi_1 \cdots \psi_r \beta([w_0])| = |w_0| + 2$, we know that $|\psi_2 \cdots \psi_r \beta([w_0])| = |w_0|$, and therefore also that $|(B, x)\psi_2 \cdots \psi_r \beta([w_0])| = |w_0|$. Then by setting $\alpha' = \gamma_1 \cdots \gamma_p$ and $\alpha' = \psi_2 \cdots \psi_r$, we are done. $\square$

Theorem 4.17 *If $\gamma \in \operatorname{Aut}^0 A_\Gamma$ with $\gamma([w_0]) = [w_0]$, then there is a factorization of γ as a product of elements of Ω that is peak-reduced with respect to $[w_0]$.*

Proof By Theorem 2.1, we factor γ as $\alpha\beta$, where $\beta \in \langle \Omega_s \rangle$ and $\alpha \in \langle \Omega_\ell^0 \rangle$. By Lemma 4.15, we write β as $\sigma' \tau_{x_1,y_1} \cdots \tau_{x_r,y_r} \iota$, where σ' is a permutation, ι is w_0- irrelevant, and the τ_{x_i,y_i} are short-range transvections such that $\{x_1, \ldots, x_r\}$ lie over distinct vertices and each x_i is domination-minimal among $\{x_i, \ldots, x_r\}$. Now we rewrite $\alpha\sigma'$ as $\sigma'\alpha'$, where $\alpha' \in \langle \Omega_\ell^0 \rangle$. By Theorem 2.1, we have a factorization $\alpha' = \sigma'' \delta_1 \cdots \delta_p$ which is peak-reduced with respect to $\tau_{x_1,y_1} \cdots \tau_{x_r,y_r} \iota([w_o])$, where σ'' is a permutation automorphism and each δ_i is a nonpermutation automorphism in Ω_ℓ^0. We set $\sigma = \sigma'\sigma''$, so that we have $\gamma = \sigma \delta_1 \cdots \delta_p \tau_{x_1,y_1} \cdots \tau_{x_r,y_r} \iota$.

Since each $x_i \in \operatorname{supp} w_0$, we deduce that no τ_{x_i,y_i} changes the support of w_0 (if it did, this would contradict Lemma 4.10). From Corollary 4.11, we know that $|\tau_{x_1,y_1} \cdots \tau_{x_r,y_r}([w_0])| \geq |w_0|$, so since α is peak-reduced, each δ_i either shortens $\delta_{i+1} \cdots \delta_p \tau_{x_1,y_1} \cdots \tau_{x_r,y_r}([w_0])$ or leaves its length unchanged. Since δ_i acts by a single multiplier, this means δ_i either leaves $\operatorname{supp} \delta_{i+1} \cdots \delta_p \tau_{x_1,y_1} \cdots \tau_{x_r,y_r}([w_0])$ the same or removes a single element. However, if this support is equal to $\operatorname{supp} w_0$, removing an element would contradict Lemma 4.10. So inductively, we deduce that

$$\operatorname{supp} \delta_1 \cdots \delta_p \tau_{x_1,y_1} \cdots \tau_{x_r,y_r}([w_0]) = \operatorname{supp} w_0.$$

Since $\delta_1 \cdots \delta_p \tau_{x_1,y_1} \cdots \tau_{x_r,y_r}([w_0])$ differs from $[w_0]$ by the permutation σ, we know that their lengths are the same.

Now inductively assume that we have written $\delta_1 \cdots \delta_p \tau_{x_1,y_1} \cdots \tau_{x_r,y_r}$ as a product $\phi_0(A_1, x_1)\phi_1 \cdots \phi_{j-2}(A_{j-1}, x_{j-1})\phi'_{j-1} \tau_{x_j,y_j} \cdots \tau_{x_r,y_r}$, with $\phi'_{j-1} \in \langle \Omega_\ell^0 \rangle$ and with $\phi_i \in \langle \Omega_\ell^0 \rangle$ and $(A_i, x_i) \in \Omega$ for each i. Also suppose that for each i,

$$(A_i, x_i)\phi_i \cdots (A_{j-1}, x_{j-1})\phi'_{j-1} \tau_{x_j,y_j} \cdots \tau_{x_r,y_r}([w_0])$$

and $$\phi_i \cdots (A_{j-1}, x_{j-1})\phi'_{j-1} \tau_{x_j,y_j} \cdots \tau_{x_r,y_r}([w_0])$$

have the same length and support as w_0. The base case for this induction has $\phi'_0 = \delta_1 \cdots \delta_p$.

Then we simply apply Lemma 4.16 to $\phi'_{j-1} \tau_{x_j,y_j} \cdots \tau_{x_r,y_r}$ and get the same statement with j instead of $j-1$. After we have done this r times, we get

$$\gamma = \sigma \phi_0(A_1, x_1)\phi_1 \cdots (A_r, x_r)\phi'_r \iota.$$

Peak-reduce each ϕ_i with respect to $(A_i, x_i)\phi_i \cdots (A_r, x_r)\phi_r([w_0])$, peak-reduce ϕ'_r with respect to $[w_0]$, and write out ι as a product of w_0-irrelevant transvections; this is a peak-reduced factorization of γ. $\square$

The following ideas appear for free groups in Lyndon–Schupp [10, Chapter I.4] and are closely related to the work of McCool in [13].

Definition 4.18 We construct a labeled, directed multigraph $\overline{\Delta}$ whose vertices are conjugacy classes of A_Γ with length equal to $|w_0| = 4k$, where we place a directed edge from $[u]$ to $[v]$ if there is a Whitehead automorphism $\alpha \in \Omega$ with $\alpha([u]) = [v]$. We label this directed edge by α. The *Whitehead graph* Δ of $[w_0]$ is the (undirected) connected component of $[w_0]$ in $\overline{\Delta}$.

Since there are only finitely many words of length $4k$, there are only finitely many conjugacy classes of length $4k$. Since Ω is finite, this means that Δ is a finite graph with only finitely many edges between any two vertices.

Corollary 4.19 *The group* $(\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$ *of automorphisms in* $\mathrm{Aut}^0 A_\Gamma$ *preserving* $[w_0]$ *is finitely generated.*

Proof A path in Δ determines an element of $\mathrm{Aut}^0 A_\Gamma$ by composing the labels along the edges. Further, if α is the automorphism determined by a path from the vertex $[w_1]$ to the vertex $[w_2]$, we know that $\alpha([w_1]) = [w_2]$ (this is true by definition for a path of length one and remains true under concatenation). In particular, this defines a map $\pi_1(\Delta, [w_0]) \to (\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$. If $\alpha \in (\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$, then by Theorem 4.17, there is a factorization $\beta_k \cdots \beta_1$ of α by elements of Ω that is peak-reduced with respect to $[w_0]$. By Corollary 4.11, this means that for each i, $\beta_i \cdots \beta_1([w_0])$ is a vertex in Δ, and β_i is an edge from $\beta_{i-1} \cdots \beta_1([w_0])$ to $\beta_i \cdots \beta_i([w_0])$. So $\beta_k \cdots \beta_1$ describes a path in Δ that maps to α. So the finitely generated group $\pi_1(\Delta, [w_0])$ surjects on $(\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$. $\qquad\qquad\square$

Remark 4.20 There are normal forms for elements of A_Γ (see Van Wyk [19], for example), so there is an effective procedure to produce the Whitehead graph of w_0. Of course, this means that there is a procedure to give a generating set for $(\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$. Unfortunately, Δ can be large for simple examples and it appears to be difficult to use this method to write down specific generating sets. McCool has explored this for the case where A_Γ is a free group and $(\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$ is a mapping class group in [14]. This procedure does not produce the familiar generating sets for the mapping class group given by Dehn twists.

4.4 Finite generation of $\mathrm{Mod}(\Gamma, w_0, Q)$

This subsection is devoted to finishing the proof of Theorem A. Recall the Whitehead graph Δ from Definition 4.18.

Lemma 4.21 *The graph* Δ *has a maximal tree* T *such that the set of all edges in* T *that are nonpermutation automorphisms forms a subtree containing* $[w_0]$.

Proof We take Δ' to be the connected component of $[w_0]$ in the subgraph of Δ gotten by deleting the edges labeled with permutation automorphisms. We take T' to be a maximal tree for Δ'.

Now, for each vertex $[w]$ of Δ not in Δ', there is a path p from $[w_0]$ to $[w]$ in Δ (paths in Δ are written in function composition order). If σ is a permutation automorphism in Ω, and α is a nonpermutation Whitehead automorphism in Ω such that $\alpha \cdot \sigma$ is a segment in p, then by Day [6, Equation (R6)], $\sigma \cdot (\sigma^{-1} \alpha \sigma)$ is another segment of length two in Δ connecting the same initial and terminal vertices. Note that since σ leaves supp w_0 invariant, we know that $\sigma^{-1} \alpha \sigma \in \Omega$. We modify p by substituting this second segment in for the first one. By repeating substitutions like this whenever possible and multiplying the permutation automorphisms together as a single permutation, we get a path p' from $[w_0]$ to $[w]$ of the form

$$\sigma_{[w]} \cdot (A_m, a_m) \cdots (A_1, a_1)$$

where $\sigma_{[w]}$ is a permutation automorphism and each $(A_i, a_i) \in \Omega$.

We already have a path in T' from $[w_0]$ to $(A_m, a_m) \cdots (A_1, a_1)([w_0])$, so we can add the edge $\sigma_{[w]}$ starting at $(A_m, a_m) \cdots (A_1, a_1)([w_0])$ to T' to get a tree containing $[w]$. It is obvious that if we add an edge gotten in this manner to T' for each vertex not in T', we will get a maximal tree for Δ satisfying the conclusions of the lemma. $\square$

Definition 4.22 For $(A, a) \in \Omega$, the *transvection set* trans(A, a) is the set of $x \in X$ with $x \in A$ or $x^{-1} \in A$, but not both.

Lemma 4.23 *The graph Δ of w_0 has a maximal tree T satisfying the following condition: for each edge α originating at a vertex $[w]$, either α is a permutation automorphism or a Whitehead automorphism $\alpha = (A, a)$ with $a \in$ supp w and* trans $\alpha \subset$ supp w.

Proof Start with an arbitrary maximal tree T_0. Suppose we have an edge (A, a) of T_0 originating at a vertex $[w]$. Fix a cyclic representative w of $[w]$, and consider the obvious representative w' of $(A, a)([w])$ based on w. If $a \notin$ supp w, then w' is the same as w with some instances of a added in. Since these are both graphically reduced representatives of conjugacy classes of the same length, we deduce that in fact, $(A, a)([w]) = [w]$. However, since T_0 is a tree, we cannot have a loop (A, a), so it must be that $a \in$ supp w. If trans$(A, a) \not\subset$ supp w, then we can rewrite (A, a) as $(A_1, a)(A_2, a)$ where trans$(A_1, a) \subset$ supp w and (A_2, a) is a product of transvections with trans$(A_2, a) \cap$ supp $w = \varnothing$. In this case, we know that $(A_2, a)([w]) = [w]$, and therefore $(A_1, a)([w]) = (A, a)([w])$. We replace the edge (A, a) with the edge (A_1, a).

Of course, we can repeat this procedure with each edge of T_0 to obtain a tree T which satisfies the conclusions of the lemma.   □

Note that if the tree T_0 above satisfies the conclusions of Lemma 4.21, then the final tree T does as well. So at this point we fix a maximal tree T in Δ that satisfies the conclusions of Lemma 4.21 and Lemma 4.23. Let T' be the subtree of T whose edges are nonpermutation Whitehead automorphisms. For each vertex $[w] \in T'$, let $\alpha_{[w]} \in \operatorname{Aut}^0 A_\Gamma$ be the product of edge labels in the edge path in T' from $[w_0]$ to $[w]$.

Lemma 4.24 *For $[w] \in T'$, we have $\operatorname{supp}[w] = \operatorname{supp} w_0$. In particular, $\alpha_{[w]}$ is an automorphism of the form $(A_m, a_m) \cdots (A_1, a_1)$ where for each i, $a_i \in \operatorname{supp} w_0$ and $\operatorname{trans}(A_i, a_i) \subset \operatorname{supp} w_0$.*

Proof Suppose $[w], [w'] \in \Delta$ and we have (A, a) with $a \in \operatorname{supp}[w]$ and $(A, a)([w]) = [w']$. Then $\operatorname{supp}[w'] \subset \operatorname{supp}[w] \cup \{a\} = \operatorname{supp}[w]$. Since the vertices of Δ are conjugacy classes of the same length as $[w_0]$, it follows that $\operatorname{supp}[w'] = \operatorname{supp}[w]$. It then follows from the definition of T' that for all $[w] \in T'$, we have $\operatorname{supp}[w] = \operatorname{supp} w_0$. The second statement in the lemma then follows from the first one.   □

Definition 4.25 For $[w] \in T'$ (possibly $[w] = [w_0]$), the automorphism $\alpha_{\beta([w])}^{-1} \beta \alpha_{[w]}$ in $(\operatorname{Aut}^0 A_\Gamma)_{[w_0]}$ is an *edge generator* if β is an edge in δ originating at $[w]$ with $\beta([w]) \in T'$ and either

- β is a permutation automorphism fixing $(\operatorname{supp} Q)^{\pm 1}$ pointwise, or
- $\beta = (B, b)$ and $\operatorname{trans} \beta \subset \operatorname{supp} w_0$.

Define the set S_e to be the set of edge generators.

Define the set $S_i \subset (\operatorname{Aut}^0 A_\Gamma)_{[w_0]}$, the set of w_0*–independent generators*, to be the set of elements $\tau_{a,b}$ where $b \in \operatorname{supp} Q$ and $a \in X$ with $a \geq b$, together with the inversions with respect to elements of $\operatorname{supp} Q$.

Define the set $S_Q \subset (\operatorname{Aut}^0 A_\Gamma)_{([w_0], Q)}$, the set of *lifted Q–transvections and Q–inversions*, to be the set of permutation automorphisms inducing a Q–inversion together with those products of transvections of length 1 or 2 that induce standard dominated Q–transvections in $\operatorname{Aut} H_\Gamma$.

Define the set $S_k \subset (\operatorname{Aut}^0 A_\Gamma)_{([w_0], Q)}$, the set of *kernel generators*, to be the set of elements of the following forms:

- automorphisms $\tau_{[x,y],c}$ (as in Section 3) where $x, y \in X$, $c \in \operatorname{supp} Q$ and $x, y \geq c$,
- partial conjugations $c_{x,\{c\}}$, where $x \in X$, $c \in \operatorname{supp} Q$ and $x \geq c$, and
- conjugations c_x, where $x \in X$.

We call S_k the set of kernel generators because these generators will be part of our generating set and they lie in the kernel of the homology representation.

Our next intermediate goal is the following.

Proposition 4.26 *The finite set $S_e \cup S_k \cup S_i$ generates $(\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$.*

We will prove some lemmas before proving this proposition.

Lemma 4.27 *Let $b \in (\mathrm{supp}\, Q)^{\pm 1}$. Suppose $\alpha = (A_m, a_m) \cdots (A_1, a_1)$ is a product of Whitehead automorphisms such that $\mathrm{v}(a_i) \neq b$ and $\mathrm{trans}(A_i, a_i) \subset \mathrm{supp}\, w_0$ for each i. Further suppose we have $\beta_1, \beta_2 \in \mathrm{Aut}^0 A_\Gamma$ both of which satisfy the following three conditions: (1) β_i fixes each $x \in X - \{b\}$; (2) $\beta_i(b)$ contains only a single instance of b and no instance of b^{-1}; and (3) for each $y \in \mathrm{supp}\, \beta_i(b)$, we have $y \geq b$ or $y = b$. Finally, suppose that $(\alpha^{-1}\beta_1\alpha\beta_2)_* \in \mathrm{Aut}\, H_\Gamma$ is the identity.*

Then we have $\alpha^{-1}\beta_1\alpha\beta_2 \in \langle S_k \rangle$.

Proof Let $\gamma = \alpha^{-1}\beta_1\alpha\beta_2$. If $x \in X - \{b\}$, then since each $\mathrm{v}(a_i) \neq b$, we know that $b \notin \mathrm{supp}\, \alpha([x])$ and therefore that γ fixes x. Since $b \notin \mathrm{trans}(A_i, a_i)$ for any i, if we alter α by an inner automorphism, we may assume that each (A_i, a_i) fixes b. Since we aim to show γ is in $\langle S_k \rangle$, which contains the inner automorphisms, we can do this. If some (A_i, a_i) conjugates some $x \geq b$ while fixing b, we can deduce that $a_i \geq b$, and it follows that each element of $\mathrm{supp}\, \gamma(b)$ is either equal to b or dominates b. Since each $a_i \neq b$, we know that only a single instance of b appears in $\gamma(b)$. Also, we know that γ fixes the image of b in H_Γ, so each element of $X - \{b\}$ that appears in $\gamma(b)$ appears in pairs of opposite exponent.

We claim that we can reduce γ to the identity by a series of applications of elements of S_k. Let the cyclic word v_0 be a graphically reduced representative of $\gamma(b)$; by the previous reasoning, we know that v_0 contains a single instance of b.

Suppose the b in v_0 is in a subsegment xby for $x, y \in L$. Note that v_0 with bxy substituted for xby represents $c_{x,\{b\}}([v_0])$ and that v_0 with xyb substituted for xby represents $c_{y,\{b\}}^{-1}([v_0])$. In this manner, by applying some partial conjugations from S_k, we can send $[v_0]$ to a conjugacy class represented by v_0 with b moved to any position in the cyclic word. If the b in v_0 is in a subsegment bxy with $\mathrm{v}(x) \neq \mathrm{v}(y)$, then v_0 with byx substituted for bxy is a representative for $\tau_{[y,x],b}([v_0])$. So by applying some elements from S_k, we can send $[v_0]$ to a conjugacy class represented by v_0 with the two letters to the right of b swapped.

So, to shorten v_0, identify an instance of some x and an instance of x^{-1} in v_0, apply elements of S_k to move b to the left of x, apply an element to move x to the right

(by a swap), move b to the right, and repeat, until x is next to x^{-1} and they cancel. Note that all of these moves are allowed since the elements appearing in v_0 other than b dominate b, and also note that these moves fix every element other than b. By this procedure we can shorten v_0 until only b remains, and we have produced an automorphism $\delta \in \langle S_k \rangle$ with $\delta = \gamma^{-1}$. $\square$

Lemma 4.28 *Let* $[w] \in T'$ *and let* $b \in \operatorname{supp} Q$. *For any* $a \in X$ *with* $a \geq b$, *we have an automorphism* γ *that is a product of transvections acting only on* b, *such that*

$$\alpha_{[w]}^{-1} \tau_{a,b} \alpha_{[w]} \gamma \in \langle S_k \rangle.$$

For any element $\alpha = \alpha_{(C,c)([w])}^{-1}(C,c)\alpha_{[w]} \in S_e$, *with* $\mathrm{v}(c) \neq b$, *there is a product* γ *of transvections acting only on* b *such that*

$$\alpha^{-1} \tau_{a,b} \alpha \gamma^{-1} \in \langle S_k \rangle.$$

If $\beta \in \langle S_k \rangle$ *fixes every element of* $X - \{b\}$, *then*

$$\alpha_{[w]}^{-1} \beta \alpha_{[w]} \in \langle S_k \rangle.$$

Proof For the first statement, note that

$$(\alpha_{[w]})_*^{-1}[a] = \sum_{i=1}^{m} p_i[c_i] \in H_\Gamma$$

for some $c_i \in X$ and nonzero integers p_i. Then for each i we have $c_i \geq a$ by Lemma 4.6, so $c_i \geq b$ and we can take $\gamma = \tau_{c_1,b}^{p_1} \cdots \tau_{c_m,b}^{p_1}$.

Similarly, it follows from Lemma 4.6 that the element γ needed for the second statement also exists.

Then the lemma is immediate from Lemma 4.24 and Lemma 4.27. $\square$

Lemma 4.29 *Let* $[w] \in T'$ *and* $\gamma \in \langle S_i \rangle$. *Then there is an element* $\gamma' \in \langle S_i \rangle$ *such that* $\alpha_{[w]}^{-1} \gamma \alpha_{[w]} \gamma'$ *is in* $\langle S_k \rangle$.

Proof We proceed by induction on the S_i–length of γ. Suppose $\gamma = \gamma_0 \beta$ where $\beta \in S_i$ and we have some $\gamma_0' \in \langle S_i \rangle$ such that $\delta = \alpha_{[w]}^{-1} \gamma_0 \alpha_{[w]} \gamma_0' \in \langle S_k \rangle$. If β is the inversion with respect to any element of $\operatorname{supp} Q$, then a computation shows that $\alpha_{[w]}$ commutes with β, and therefore

$$\alpha_{[w]}^{-1} \gamma \alpha_{[w]} \beta^{-1} \gamma_0' = \delta \in \langle S_k \rangle$$

and we can take $\gamma' = \beta^{-1} \gamma_0'$.

If β is a transvection $\tau_{a,b}$ with $b \in \operatorname{supp} Q$, then by Lemma 4.28 we have a $\gamma'' \in \langle S_i \rangle$ with $\delta' = \alpha_{[w]}^{-1} \beta \alpha_{[w]} \gamma'' \in \langle S_k \rangle$. Set $\delta'' = (\gamma_0')^{-1} \delta' \gamma_0'$. From Lemma 3.4 (with $Z = \operatorname{supp} Q$), we have $\delta'' \in \langle S_k \rangle$. We set $\gamma' = \gamma'' \gamma_0'$, and we have

$$\alpha_{[w]}^{-1} \gamma_0 \beta \alpha_{[w]} \gamma'' \gamma_0' = \alpha_{[w]}^{-1} \gamma_0 \alpha_{[w]} \delta' \gamma_0'$$
$$= \delta (\gamma_0')^{-1} \delta' \gamma_0' = \delta \delta''$$

which is in $\langle S_k \rangle$. □

Proof of Proposition 4.26 First note that S_e is finite because Δ is finite, and S_i and S_k are finite because X is finite. Now suppose we have an edge α between two vertices $[w_1]$ and $[w_2]$ of Δ. We know that $[w_i] = \sigma_i([v_i])$ where σ_i is a possibly trivial permutation automorphism and $[v_i] \in T'$, for $i = 1, 2$. Define the set $S \subset \operatorname{Aut} A_\Gamma$ to be the set of elements of the form $\alpha_{[v_2]}^{-1} \sigma_2^{-1} \alpha \sigma_1 \alpha_{[v_1]}$, indexed over all edges α of Δ. Since T' is a maximal tree for Δ, the elements of S describe a generating set for $\pi_1(\Delta, [w_0])$, and as explained in Corollary 4.19, they therefore generate $(\operatorname{Aut}^0 A_\Gamma)_{[w_0]}$.

Since it is obvious that $S_e \cup S_k \cup S_i \subset (\operatorname{Aut}^0 A_\Gamma)_{[w_0]}$, we prove the lemma by showing that $S \subset \langle S_e \cup S_k \cup S_i \rangle$. Consider an arbitrary element of S:

$$\beta = \alpha_{[v_2]}^{-1} \sigma_2^{-1} \alpha \sigma_1 \alpha_{[v_1]}.$$

If α is a permutation automorphism, then we write $\sigma_2^{-1} \alpha \sigma_1$ as a single permutation σ_3. Since $[v_1]$ and $[v_2]$ are both in T', we know that $\operatorname{supp}[v_1] = \operatorname{supp}[v_2] = \operatorname{supp} w_0$, and therefore σ_3 leaves $\operatorname{supp} w_0$ invariant. It follows from this and the fact that σ_3 is in $\operatorname{Aut}^0 A_\Gamma$ that σ_3 factors as a product of a permutation automorphism $\sigma_{w_0} \in \operatorname{Aut}^0 A_\Gamma$ that fixes $(\operatorname{supp} Q)^{\pm 1}$ pointwise and a permutation automorphism $\sigma_Q \in \operatorname{Aut}^0 A_\Gamma$ that fixes $(\operatorname{supp} w_0)^{\pm 1}$ pointwise. Since σ_Q is in $\operatorname{Aut}^0 A_\Gamma$ and fixes $(\operatorname{supp} w_0)^{\pm 1}$, it follows that $\sigma_Q \in \langle S_i \rangle$. We know $\operatorname{supp}[v_1] = \operatorname{supp} w_0$, so σ_Q fixes $[v_1]$ and therefore $\sigma_{w_0}([v_1]) = [v_2]$ and $\alpha_{[v_2]}^{-1} \sigma_{w_0} \alpha_{[v_1]} \in S_e$. Then β will be in $\langle S_e \cup S_k \rangle$ if $\alpha_{[v_1]}^{-1} \sigma_Q \alpha_{[v_1]}$ is. Since $\sigma_Q \in \langle S_i \rangle$, Lemma 4.29 says that there is a $\gamma \in \langle S_i \rangle$ with $\alpha_{[v_1]}^{-1} \sigma_Q \alpha_{[v_1]} \gamma \in \langle S_k \rangle$. The proposition follows in this case.

So assume that $\alpha = (A, a)$. By replacing α with $\sigma_2^{-1} \alpha \sigma_2$ and σ_1 with $\sigma_2^{-1} \sigma_1$, we may assume that $\sigma_2 = 1$. If we set $[w_3] = \alpha^{-1} \alpha_{[v_2]}([w_0])$, then it follows from the construction of T' that $[w_3] \in T'$. Then we know that $\alpha_{[w_3]}^{-1} \sigma_1 \alpha_{[v_1]}$ is in S_e, and therefore β is in $\langle S_e \cup S_k \cup S_i \rangle$ only if the element

$$\beta' = \alpha_{[v_2]}^{-1} \alpha \alpha_{[w_3]}$$

is as well. We may rewrite $\alpha = (A', a) \gamma$, where $\gamma \in \langle S_i \rangle$ and $\operatorname{trans}(A', a) \subset \operatorname{supp} w_0$. Since $[w_3] \in T'$, we know $\operatorname{supp}[w_3] = \operatorname{supp}[w_0]$, and γ fixes $[w_3]$. We may rewrite β'

as the product of the element $\alpha_{[v_2]}^{-1}(A', a)\alpha_{[w_3]}$, which is in S_e, with the element $\alpha_{[w_3]}^{-1}\gamma\alpha_{[w_3]}$. By Lemma 4.29, there is an element $\gamma' \in \langle S_i \rangle$ such that $\alpha_{[w_3]}^{-1}\gamma\alpha_{[w_3]}\gamma'$ is in $\langle S_k \rangle$. The proposition follows. $\qquad\square$

We proceed by showing that we can do better:

Proposition 4.30 *Any element of* $(\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$ *can be written as the product of a single element of* $\langle S_e \cup S_k \rangle$ *and a single element of* $\langle S_i \rangle$.

Sublemma 4.31 *Suppose* (B, b) *is a Whitehead automorphism,* $a \in L$ *with* $a \geq b$ *and* $a, a^{-1} \notin B$. *Suppose* $[u]$ *is an element or conjugacy class in* A_Γ *with* $\mathrm{v}(b) \notin \mathrm{supp}[u]$ *and* v *is a graphically reduced word or cyclic word representing* $(B, b)([u])$. *Then if* v' *is* v *with all instances of* b *replaced by* a *and all instances of* b^{-1} *replaced by* a^{-1}, *then* v' *is a representative of* $(B - b + a, a)([u])$.

Proof Note that since $a \geq b$, $(B - b + a, a)$ is well defined by Day [6, Lemma 2.5]. Pick a graphically reduced representative u for $[u]$; obtain a representative $\tilde{v}$ for $(B, b)([u])$ by applying (B, b) letter-by-letter to u. It is immediate that if $\tilde{v}'$ is $\tilde{v}$ with these substitutions, then $\tilde{v}$ represents $(B - b + a, a)$. Note that $\mathrm{v}(b) \notin \mathrm{supp}\,\tilde{v}$. Since a commutes with every letter that b commutes with (except possibly $b^{\pm 1}$), each time we modify $\tilde{v}$ by swapping two adjacent, commuting letters, or by making a graphic reduction, we can make a parallel modification to $\tilde{v}'$ and still have representatives that differ by the described substitution and represent the same two elements. Since we can get from any representative of $(B, b)([u])$ to the representative v by such moves, we have proven the statement. $\qquad\square$

Lemma 4.32 *Suppose* $\beta = (B, b) \in \Omega$ *such that* $\mathrm{v}(b) \notin \mathrm{supp}\, w_0$ *and for some* $[w] \in T'$, $\alpha_{[w]}^{-1}\beta\alpha_{[w]} \in S_e$. *If* $a \in (\mathrm{supp}\, w_0) \cap (\mathrm{trans}\, \beta)$, *then* $a \not\sim b$.

Proof Suppose for contradiction that $a \sim b$. By the construction of T', w has the same length and support as w_0, so by Lemma 4.10, there is a single instance of a and a single instance of a^{-1} in w. So write w as the graphically reduced cyclic word $aua^{-1}v$. Then $[u, a] \neq 1$ and $[v, a] \neq 1$. Since $a \sim b$, we know that $[u, b] \neq 1$ and $[v, b] \neq 1$ as well. Since $a \in \mathrm{trans}\,\beta$, we may assume that $a \in B$ and $a^{-1} \notin B$ (the case where $a^{-1} \in B$ and $a \notin B$ is similar). We know $\beta([w]) = [w]$ and $\beta(a) = ab$. Since $[v, b] \neq 1$, β cannot send v to an element represented by a reduced word ending in b^{-1} or beginning with b. It follows that $\beta(u) = b^{-1}ub$ and $\beta(v) = v$.

Since $a \sim b$, we know from Day [6, Lemma 2.5] that $(B - b, a)$ is a well-defined Whitehead automorphism. We know b does not appear in u or v, so by Sublemma 4.31

$(B - b, a)(u) = a^{-1}ua$ and $(B - b, a)(v) = v$. Since $[u, a] \neq 1$, these expressions are graphically reduced. Then $(B - b, a)(w) = aa^{-1}uaa^{-1}v = uv$, and $|uv| \leq |w| - 2$. This contradicts Corollary 4.11. $\qquad\square$

Lemma 4.33 *Suppose $(B, b) \in \Omega$ such that $\mathrm{v}(b) \notin \mathrm{supp}\, w_0$ and for some $[w] \in T'$, $\alpha_{[w]}^{-1}\beta\alpha_{[w]} \in S_e$. Suppose $a \in L$ such that $a \geq b$ and $a, a^{-1} \notin B$. Then $(B - b + a, a)$ fixes $[w]$, and $\alpha_{[w]}^{-1}(B - b + a, a)\alpha_{[w]} \in S_e$.*

Proof It is immediate from Sublemma 4.31 that $(B - b + a, a)$ fixes $[w]$. Note that $\mathrm{trans}(B - b + a, a) = \mathrm{trans}(B, b) \subset \mathrm{supp}\, w_0$, so $\alpha_{[w]}^{-1}(B - b + a, a)\alpha_{[w]} \in S_e$. $\qquad\square$

Lemma 4.34 *Suppose $\alpha_{[w]}^{-1}(B, b)\alpha_{[w]} \in S_e$ and $\mathrm{v}(b) \in \mathrm{supp}\, Q$. If γ is a product of transvections acting only on b then $\gamma\alpha_{[w]}^{-1}(B, b)\alpha_{[w]}\gamma^{-1}$ is in $\langle S_e \cup S_k \rangle$.*

Proof Let $S_{[w]}$ be the union of S_k with the set of $(C, c) \in \Omega$ such that $(C, c)([w]) = [w]$ and $\mathrm{trans}(C, c) \subset \mathrm{supp}\, w_0$.

As a base case, consider the effect of a single transvection $\tau_{a,b}$ on a (B, b), where $a \in L$ and $a \geq b$. Note that if $a \in \mathrm{trans}(B, b)$, then $a \in \mathrm{supp}\, w_0$ and $b \geq a$. Then $a \sim b$, contradicting Lemma 4.32. So we know $a \notin \mathrm{trans}(B, b)$. Possibly by multiplying (B, b) by an inner automorphism, we may assume $a \notin B$. Then since $a \notin \mathrm{trans}(B, b)$, we know $a^{-1} \notin B$. Since $\tau_{a,b} = (\{a, b\}, a)$ and $b^{-1} \notin \{a, b\}$, we have the following special case of Equation (R4) of Day [6]:

$$(B, b)^{-1}\tau_{a,b}(B, b) = \tau_{a,b}(B - b + a, a).$$

We rephrase this as

$$\tau_{a,b}(B, b)\tau_{a,b}^{-1} = (B, b)(B - b + a, a).$$

Similarly, note that

$$\tau_{a,b^{-1}}(B, b)\tau_{a,b^{-1}}^{-1} = (B - b + a^{-1}, a^{-1})(B, b).$$

By Lemma 4.33, $(B - b + a, a)$ and $(B - b + a^{-1}, a^{-1})$ both preserve $[w]$. Of course, $\mathrm{trans}(B - b + a^{-1}, a^{-1}) = \mathrm{trans}(B - b + a, a) = \mathrm{trans}(B, b) \subset \mathrm{supp}\, w_0$.

Now consider $\tau_{a',b}$ for some $a' \in L$ with $a \geq b$. If $\mathrm{v}(a') = \mathrm{v}(a)$, then $\tau_{a,b}$ and $(B - b + a, a)$ commute. Since $\mathrm{trans}(B - b + a, a) = \mathrm{trans}(B, b)$, and since by Lemma 4.32, we know $a' \notin \mathrm{trans}(B, b)$, we know $a' \notin \mathrm{trans}(B - b + a, a)$. Further, we know that $b, b^{-1} \notin B - b + a$. Then by Day [6, Equation (R3)], $\tau_{a',b}$ and $(B - b + a, a)$ commute up to an inner automorphism,. Similarly, $\tau_{a',b^{-1}}$ and $(B - b + a, a)$ commute up to an inner automorphism.

So if γ is a product of transvections acting only on b, then up to inner automorphisms, $\gamma(B,b)\gamma^{-1}$ is (B,b) times some number of elements of the form $(B-b+a,a)$ for various $a \geq b$. Then in particular, $\gamma(B,b)\gamma^{-1}$ is in $\langle S_{[w]} \rangle$.

Now suppose γ is a product of transvections acting only on b. From Lemma 4.28, there is a product γ' of transvections acting on b and a $\delta \in \langle S_e \rangle$ with $\gamma\alpha_{[w]} = \delta\alpha_{[w]}^{-1}\gamma'$. We deduce from the previous paragraph there is an element $\beta \in \langle S_{[w]} \rangle$ with $\gamma'(B,b)\gamma'^{-1} = \beta$. Then we have

$$\gamma\alpha_{[w]}^{-1}(B,b)\alpha_{[w]}\gamma^{-1} = \delta\alpha_{[w]}^{-1}\beta\alpha_{[w]}\delta^{-1}.$$

By Lemma 4.28, if $\delta' \in S_k$, then $\alpha_{[w]}^{-1}\delta'\alpha_{[w]} \in \langle S_k \rangle$. And if $(C,c) \in S_{[w]}$, then $\alpha_{[w]}^{-1}(C,c)\alpha_{[w]} \in S_e$. So $\alpha_{[w]}^{-1}\beta\alpha_{[w]} \in \langle S_e \cup S_k \rangle$, proving the lemma. $\square$

Lemma 4.35 Suppose $b \in \mathrm{supp}\, Q$, γ_1 is a product of transvections acting on b and $\alpha \in S_e \cup S_k$. Then there is a $b' \in \mathrm{supp}\, Q$ and a product γ_2 of transvections acting on b' such that $\gamma_1\alpha\gamma_2^{-1} \in \langle S_e \cup S_k \rangle$.

Proof First we note that if $\alpha \in S_k$, then we can take $\gamma_2 = \gamma_1$ and the lemma follows from the identities in Sublemma 3.2 and Sublemma 3.3.

Now suppose $\alpha = \alpha_{\beta([w])}^{-1}\beta\alpha_{[w]}$ and β is a permutation automorphism. Since $[w]$ and $\beta([w])$ are both in T', we know that $\mathrm{supp}[w] = \mathrm{supp}\,\beta([w]) = \mathrm{supp}\, w_0$. So β leaves $\mathrm{supp}\, Q$ invariant and we can set $b' = \mathrm{v}(\beta^{-1}(b)) \in \mathrm{supp}\, Q$. By Lemma 4.6, we can find a product γ_2 of transvections acting only on b', such that γ_2 sends the image of b' in H_Γ to the same element that $\alpha_{[w]}^{-1}\beta^{-1}\alpha_{\beta[w]}\gamma_1^{-1}$ sends it to. Let $\alpha' = \beta\alpha_{[w]}\beta^{-1}$ and let $\gamma' = \beta\gamma_2\beta^{-1}$. Then by Day [6, Equation (R6)], α' can be written as a product of nonpermutation automorphisms with multipliers not equal to $\mathrm{v}(b)$ and γ' can be written as a product of transvections acting only on b. It then follows from Lemma 4.27 that

$$\gamma_1(\alpha_{\beta([w])}^{-1}\beta\alpha_{[w]})\gamma_2(\alpha_{[w]}^{-1}\beta^{-1}\alpha_{\beta([w])}) = \gamma_1(\alpha_{[w]}\alpha')\gamma'(\alpha_{[w]}\alpha')^{-1} \in \langle S_k \rangle$$

which proves the lemma in this case.

If $\alpha = \alpha_{(C,c)([w])}^{-1}(C,c)\alpha_{[w]}$, then the lemma follows from Lemma 4.34 if $\mathrm{v}(c) = \mathrm{v}(b)$ and from Lemma 4.28 if $\mathrm{v}(c) \neq \mathrm{v}(b)$. $\square$

Proof of Proposition 4.30 For each $b \in \mathrm{supp}\, Q$, take $S(b)$ to be the subgroup generated by $\{\tau_{a,b} | a \in X, a \geq b\}$ and the inversion with respect to b, and take:

$$S = \bigcup_{b \in \mathrm{supp}\, Q} S(b)$$

For $\alpha \in (\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$, take $d(\alpha)$ to be the minimum number of elements of S appearing in any factorization of α as a product of elements of $(S_k \cup S_e \cup S)^{\pm 1}$. Note $S_i \subset S$, so such a factorization exists by Proposition 4.26.

We will prove this proposition by induction on $d(\alpha)$. If $d(\alpha) = 0$, then the proposition is obviously true. Now suppose that we have $\alpha = \beta \gamma \alpha'$, where $\gamma \in S(b)$ for some $b \in \mathrm{supp}\, Q$, $\beta \in \langle S_e \cup S_k \rangle$ and $\alpha' \in (\mathrm{Aut}^0 A_\Gamma)_{[w_0]}$ with $d(\alpha') = d(\alpha) - 1$. By repeated application of Lemma 4.35, we know we have some $b' \in \mathrm{supp}\, Q$, some $\gamma' \in S(b')$ and some $\beta' \in \langle S_k \cup S_e \rangle$ with $\beta \gamma = \gamma' \beta'$. Then $\alpha = \gamma' \beta' \alpha'$. Since $d(\beta' \alpha') \leq d(\alpha) - 1$, we can apply the inductive hypothesis and get $\gamma'' \in \langle S_i \rangle$ and $\beta'' \in \langle S_k \cup S_e \rangle$ with $\beta' \alpha' = \gamma'' \beta''$. Then $\alpha = \gamma' \gamma'' \beta''$; since $\gamma' \gamma'' \in \langle S_i \rangle$, we are done. $\qquad \square$

Lemma 4.36 *We have $\langle S_k \cup S_e \rangle < (\mathrm{Aut}^0 A_\Gamma)_Q$.*

Proof For $\delta \in S_k$, since $\delta \in \ker(\mathrm{Aut}^0 A_\Gamma \to \mathrm{Aut}\, H_\Gamma)$, it is obvious that $\delta_* Q = Q$.

For $[w] \in T'$, note that $(\alpha_{[w]})_* Q = Q$. This is because $\alpha_{[w]}$ is a product of elements (A, a) with $\mathrm{trans}(A, a) \subset \mathrm{supp}\, w_0$. If $\alpha = \alpha_{\beta([w])}^{-1} \beta \alpha_{[w]} \in S_e$, then either β is a permutation fixing $(\mathrm{supp}\, Q)^{\pm 1}$ or β is a nonpermutation Whitehead automorphism with $\mathrm{trans}\, \beta \subset \mathrm{supp}\, w_0$. In either case, it follows that $\alpha_* Q = Q$. $\qquad \square$

Theorem 4.37 *The group $(\mathrm{Aut}^0 A_\Gamma)_{([w_0], Q)}$ is generated by the finite set $S_e \cup S_k \cup S_Q$.*

Proof As previously noted, S_e and S_k are finite. The set S_Q is finite because Q is finite. Now suppose that $\alpha \in (\mathrm{Aut}^0 A_\Gamma)_{([w_0], Q)}$. By Proposition 4.30, we can rewrite α as $\beta \gamma$ where $\beta \in \langle S_e \cup S_k \rangle$ and $\gamma \in \langle S_i \rangle$.

By Lemma 4.36, we know that $\beta_* Q = Q$. Since $\alpha_* Q = Q$, it follows that $\gamma_* Q = Q$. So there is an element $\delta \in \langle S_Q \rangle$ such that $\gamma \delta^{-1} \in \ker(\mathrm{Aut}^0 A_\Gamma \to \mathrm{Aut}\, H_\Gamma)$, by Theorem 4.4. Then by Proposition 3.5 (with $Z = \mathrm{supp}\, Q$), we know that $\gamma \delta^{-1} \in \langle S_k \rangle$. Since $\alpha = \beta(\gamma \delta^{-1})\delta$, we have proven the theorem. $\qquad \square$

Proposition 4.38 *The group $(\mathrm{Aut}\, A_\Gamma)_{([w_0], Q)}$ is finitely generated.*

Proof Recall that $\mathrm{Aut}^0 A_\Gamma$ is a finite-index normal subgroup of $\mathrm{Aut}\, A_\Gamma$. Then $\mathrm{Aut}^0 A_\Gamma$ is also finite-index and normal in $\langle \mathrm{Aut}^0 A_\Gamma, (\mathrm{Aut}\, A_\Gamma)_{([w_0], Q)} \rangle$. By the classical second isomorphism theorem, we have:

$$\langle \mathrm{Aut}^0 A_\Gamma, (\mathrm{Aut}\, A_\Gamma)_{([w_0], Q)} \rangle / \mathrm{Aut}^0 A_\Gamma$$

$$\cong (\mathrm{Aut}\, A_\Gamma)_{([w_0], Q)} / \big((\mathrm{Aut}\, A_\Gamma)_{([w_0], Q)} \cap \mathrm{Aut}^0 A_\Gamma\big)$$

But $(\mathrm{Aut}\, A_\Gamma)_{([w_0], Q)} \cap \mathrm{Aut}^0 A_\Gamma = (\mathrm{Aut}^0 A_\Gamma)_{([w_0], Q)}$, so $(\mathrm{Aut}^0 A_\Gamma)_{([w_0], Q)}$ is finite-index in $(\mathrm{Aut}\, A_\Gamma)_{([w_0], Q)}$. So we are done by Theorem 4.37. $\qquad \square$

Final step in the proof of Theorem A   Let $Z_{A_\Gamma}(w_0)$ denote the centralizer of w_0 in A_Γ. Consider the following sequence of maps, which we will show to be exact:

$$Z_{A_\Gamma}(w_0) \to \mathrm{Mod}(\Gamma, w_0, Q) \xrightarrow{\pi} (\mathrm{Out}\, A_\Gamma)_{([w_0], Q)} \to 0,$$

where the first map sends an element to its corresponding inner automorphism.

If $[\alpha] \in (\mathrm{Out}\, A_\Gamma)_{([w_0], Q)}$ and $\alpha \in \mathrm{Aut}\, A_\Gamma$ is a lift of α, then α sends w_0 to a conjugate $u^{-1} w_0 u$. If we compose α with the inner automorphism given by conjugation by u^{-1}, we get an automorphism in $\mathrm{Mod}(\Gamma, w_0, Q)$ that projects to $[\alpha]$. This explains the surjectivity of π.

If α is in the kernel of π, then it is the inner automorphism c_u for some $u \in A_\Gamma$. Of course, $c_u \in \mathrm{Mod}(\Gamma, w_0, Q)$ if and only if $c_u \in (\mathrm{Aut}\, A_\Gamma)_{w_0}$, which is true if and only if $u \in Z_{A_\Gamma}(w_0)$, proving the exactness of this sequence.

From Proposition 4.38, the group $(\mathrm{Out}\, A_\Gamma)_{([w_0], Q)}$ is finitely generated. Servatius's centralizer theorem from [18] completely describes the centralizers of elements in A_Γ; in particular, it tells us that $Z_{A_\Gamma}(w_0)$ is finitely generated. Since $\mathrm{Mod}(\Gamma, w_0, Q)$ surjects onto a finitely generated group with finitely generated kernel, it is finitely generated.   □

5   Closing remarks

The work in this paper opens the way for further study of mapping class groups over graphs. First of all, it would be interesting to recover Definition 1.1 by means of a geometric construction. In the extreme cases, $\mathrm{Sp}(2g, \mathbb{Z})$ can be seen as the linear automorphisms of the torus $\mathbb{T}^{2g}$ that preserve a standard symplectic differential form, and $\mathrm{Mod}_{g,1}$ can be seen as the homotopy group of self-homotopy-equivalences of a graph that preserve some additional combinatorial structure called a "fat graph" structure (see Penner [15]). It is worth noting that $\mathbb{T}^{2g}$ and certain graphs are examples of *Salvetti complexes*. The Salvetti complex S_Γ is a finite cubical complex that forms a natural $K(A_\Gamma, 1)$ space (see Charney [3, Section 2.6]). If we take the monoid of self-homotopy-equivalences of S_Γ and take a quotient by considering maps equivalent if they are homotopic, we get a group. Call this group G. Note that $G \cong \mathrm{Aut}\, A_\Gamma$. This brings us to the following problem:

Open Problem 5.1   Produce a structure on S_Γ and a corresponding symplectic structure (w, Q) on A_Γ such that the subgroup of G of elements represented by maps fixing this structure is naturally isomorphic to $\mathrm{Mod}(\Gamma, w, Q)$.

By a structure on S_Γ, I mean some extra combinatorial data, or some extra differential data, or some combination of the two.

Theorem A could be a starting point for future homological finiteness results about $\mathrm{Mod}(\Gamma, w, Q)$. This conjecture could be a possible next step.

Conjecture 5.2 *For every graph Γ with a symplectic structure (w, Q) on A_Γ, the group $\mathrm{Mod}(\Gamma, w, Q)$ is finitely presented.*

There are combinatorial methods to show that $\mathrm{Mod}_{g,1}$ is finitely presented (see Mc-Cool [13]) which could potentially be extended to prove Conjecture 5.2.

In the extreme cases, it is known that both $\mathrm{Mod}_{g,1}$ and $\mathrm{Sp}(2g, \mathbb{Z})$ contain finite index subgroups with finite $K(\pi, 1)$ complexes. This implies that both groups are of type *VFL*, which is a strong homological finiteness condition (see Brown [2, Chapter VIII.11]). This leads us to the following conjecture.

Conjecture 5.3 *For every graph Γ with symplectic structure (w, Q) on A_Γ, the group $\mathrm{Mod}(\Gamma, w, Q)$ has a finite-index subgroup G with a finite $K(G, 1)$ complex.*

It seems unlikely that Conjecture 5.3 could be proven by purely combinatorial methods, but given a solution to Open Problem 5.1, it is conceivable that one could recover such a $K(G, 1)$ complex as a kind of moduli space of Salvetti complexes with symplectic structures. A related problem would then be to find bounds on the virtual cohomological dimension of $\mathrm{Mod}(\Gamma, w, Q)$. Charney–Crisp–Vogtmann [4] and Charney–Vogtmann [5] have already made much progress on the parallel problem for $\mathrm{Aut}\, A_\Gamma$.

Theorem B could be a starting point for work on the homological properties of $\mathrm{IAut}\, A_\Gamma$. In the usual way (as with $\mathrm{Mod}_{g,1}$ or with $\mathrm{Aut}\, F_n$), the action of $\mathrm{Aut}\, A_\Gamma$ on the 2–step nilpotent truncation of A_Γ defines an $\mathrm{Aut}\, A_\Gamma$–equivariant homomorphism (a *Johnson homomorphism*) from $\mathrm{IAut}\, A_\Gamma$ to an abelian group. We can then ask the following question.

Question 5.4 For arbitrary Γ, is the image of the Johnson homomorphism on $\mathrm{IAut}\, A_\Gamma$ equal to the abelianization of $\mathrm{IAut}\, A_\Gamma$?

This question was answered in the affirmative for IA_n, independently by Cohen–Pakianathan, by Farb and by Kawazumi [8] (see Theorem 1.1 of Pettet [16]).

The following conjecture is linked to Charney–Vogtmann [5].

Conjecture 5.5 *For every graph Γ, the group $\mathrm{IAut}\, A_\Gamma$ is torsion-free and there is a finite-dimensional $K(\mathrm{IAut}\, A_\Gamma, 1)$ complex.*

A related problem is to bound the dimension of such a complex, as Bestvina–Bux–Margalit [1] did in the case of IA_n. We do not expect such a complex to have finitely many cells in each dimension, but only that such a complex would be finite-dimensional.

References

[1] **M Bestvina**, **K-U Bux**, **D Margalit**, *Dimension of the Torelli group for* Out(F_n), Invent. Math. 170 (2007) 1–32 MR2336078

[2] **K S Brown**, *Cohomology of groups*, Graduate Texts in Math. 87, Springer, New York (1982) MR672956

[3] **R Charney**, *An introduction to right-angled Artin groups*, Geom. Dedicata 125 (2007) 141–158 MR2322545

[4] **R Charney**, **J Crisp**, **K Vogtmann**, *Automorphisms of* 2–*dimensional right-angled Artin groups*, Geom. Topol. 11 (2007) 2227–2264 MR2372847

[5] **R Charney**, **K Vogtmann**, *Finiteness properties of automorphism groups of right-angled Artin groups*, to appear in Bull. London. Math. Soc. Available at `http://people.brandeis.edu/~charney/webpubs.htm`

[6] **M B Day**, *Peak reduction and presentations for automorphism groups of right-angled Artin groups*, Geom. Topol. 13 (2009) 817–855

[7] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Book draft Available at `http://www.math.utah.edu/~margalit/primer/`

[8] **N Kawazumi**, *Cohomological aspects of Magnus expansions* `arXiv: math.GT/0505497`

[9] **M R Laurence**, *A generating set for the automorphism group of a graph group*, J. London Math. Soc. (2) 52 (1995) 318–334 MR1356145

[10] **R C Lyndon**, **P E Schupp**, *Combinatorial group theory*, Classics in Math., Springer, Berlin (2001) MR1812024Reprint of the 1977 edition

[11] **W Magnus**, *Über n–dimensionale Gittertransformationen*, Acta Math. 64 (1935) 353–367 MR1555401

[12] **W Magnus**, **A Karrass**, **D Solitar**, *Combinatorial group theory. Presentations of groups in terms of generators and relations*, second edition, Dover, Mineola, NY (2004) MR2109550

[13] **J McCool**, *Some finitely presented subgroups of the automorphism group of a free group*, J. Algebra 35 (1975) 205–213 MR0396764

[14] **J McCool**, *Generating the mapping class group (an algebraic approach)*, Publ. Mat. 40 (1996) 457–468 MR1425631

[15] **R C Penner**, *The moduli space of a punctured surface and perturbative series*, Bull. Amer. Math. Soc. (N.S.) 15 (1986) 73–77 MR838790

[16] **A Pettet**, *The Johnson homomorphism and the second cohomology of* IA$_n$, Algebr. Geom. Topol. 5 (2005) 725–740 MR2153110

[17]   **J-P Serre**, *Lie algebras and Lie groups*, Lecture Notes in Math. 1500, Springer, Berlin (2006)   MR21796911964 lectures given at Harvard University, Corrected fifth printing of the second (1992) edition

[18]   **H Servatius**, *Automorphisms of graph groups*, J. Algebra 126 (1989) 34–60 MR1023285

[19]   **L VanWyk**, *Graph groups are biautomatic*, J. Pure Appl. Algebra 94 (1994) 341–352 MR1285550

[20]   **H Zieschang**, *Über Automorphismen ebener diskontinuierlicher Gruppen*, Math. Ann. 166 (1966) 148–167   MR0201521

Department of Mathematics, California Institute of Technology
1200 E California Blvd, Pasadena, CA 91101, USA

`mattday@caltech.edu`

`http://www.its.caltech.edu/~mattday`