Palindromic automorphisms of free groups and rigidity of automorphism groups of right-angled Artin groups

by

Neil James Fullarton

A thesis submitted to the
College of Science and Engineering
at the University of Glasgow
for the degree of
Doctor of Philosophy

June 2014

For Jim and June

Abstract Let F_n denote the free group of rank n with free basis X. The *palindromic automorphism group* ΠA_n of F_n consists of automorphisms taking each member of X to a palindrome: that is, a word on $X^{\pm 1}$ that reads the same backwards as forwards. We obtain finite generating sets for certain stabiliser subgroups of ΠA_n. We use these generating sets to find an infinite generating set for the so-called *palindromic Torelli group* $\mathcal{PI}_n$, the subgroup of ΠA_n consisting of palindromic automorphisms inducing the identity on the abelianisation of F_n. Two crucial tools for finding this generating set are a new simplicial complex, the so-called *complex of partial* π*-bases*, on which ΠA_n acts, and a Birman exact sequence for ΠA_n, which allows us to induct on n.

We also obtain a rigidity result for automorphism groups of right-angled Artin groups. Let Γ be a finite simplicial graph, defining the right-angled Artin group A_Γ. We show that as A_Γ ranges over all right-angled Artin groups, the order of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ does not have a uniform upper bound. This is in contrast with extremal cases when A_Γ is free or free abelian: in this case, $|\mathrm{Out}(\mathrm{Aut}(A_\Gamma))| \leq 4$. We prove that no uniform upper bound exists in general by placing constraints on the graph Γ that yield tractable decompositions of $\mathrm{Aut}(A_\Gamma)$. These decompositions allow us to construct explicit members of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$.

Acknowledgements First and foremost, I would like to thank my supervisor, Tara Brendle, for her constant support, both mathematically and personally, over the last four years. I am grateful to the Engineering and Physical Sciences Research Council for the funding with which I was provided to complete my PhD. I am also indebted to the University of Glasgow's School of Mathematics and Statistics and College of Science and Engineering for providing me with many excellent learning and teaching opportunities over the years. I am grateful to Alessandra Iozzi and the Institute for Mathematical Research at Eidgenössische Technische Hochschule Zürich, where part of this work was completed. I also wish to thank Ruth Charney, Dan Margalit, Andrew Putman and Karen Vogtmann for helpful conversations.

A debt of gratitude is owed to my parents and sister, without whom I would not be where I am today. I am also grateful to my officemates, Liam Dickson and Pouya Adrom, for their spirit of camaraderie. I would also like to thank the philosophers Anne, Thom, Luke and Rosie for all their moral support.

To Laura, I am eternally grateful for always being there and for getting me back in the habit.

And finally, to Finlay. Thanks for all the sandwiches.

I declare that, except where explicit reference is made to the contribution of others, this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Neil J. Fullarton

Contents

Chapter 1

Introduction

The goal of this thesis is to investigate the structure of certain automorphism groups of free groups and, more generally, of right-angled Artin groups. In particular, we will find explicit generating sets for certain subgroups of the so-called palindromic automorphism group of a free group, using geometric methods, as well as investigating the structure of the outer automorphism group of the automorphism group of a right-angled Artin group.

The Torelli group. Let F_n be the free group of rank n on some fixed free basis $X = \{x_1, \ldots, x_n\}$. Both F_n and its automorphism group $\mathrm{Aut}(F_n)$ are fundamental objects of study in group theory, due to the ubiquity of F_n throughout mathematics. For instance, free groups appear as fundamental groups of graphs and oriented surfaces with boundary, and every finitely generated group is the quotient of some finite rank free group. While $\mathrm{Aut}(F_n)$ has been studied for a century, there is still much to be learned about its structure. It has a rich subgroup structure, containing certain mapping class groups [29] and braid groups [7], for example. One particularly interesting subgroup of $\mathrm{Aut}(F_n)$ is IA_n, the kernel of the canonical surjection $\Psi : \mathrm{Aut}(F_n) \to \mathrm{GL}(n, \mathbb{Z})$ induced by abelianising F_n. This kernel is called the *Torelli group*, and we have the short exact sequence

$$1 \longrightarrow \mathrm{IA}_n \longrightarrow \mathrm{Aut}(F_n) \longrightarrow \mathrm{GL}(n, \mathbb{Z}) \longrightarrow 1.$$

While Magnus gave a finite generating set for IA_n in 1935 [43], it is still unknown whether IA_n is finitely presentable for $n \geq 4$ (while $\mathrm{IA}_2 \cong F_2$ [49] and IA_3 is not finitely presentable [40]).

Palindromic automorphisms of free groups. The subgroup of $\mathrm{Aut}(F_n)$ we shall study in Chapter 2 is the *palindromic automorphism group of* F_n, denoted $\Pi\mathrm{A}_n$. Introduced by Collins [18], $\Pi\mathrm{A}_n$ consists of automorphisms of F_n that send each $x \in X$ to a *palindrome*, that is, a word on $X^{\pm 1}$ that may be read the same backwards as forwards. Collins gave a finite presentation for $\Pi\mathrm{A}_n$, and it can be shown that a certain subgroup $\mathrm{P}\Pi\mathrm{A}_n \leq \Pi\mathrm{A}_n$, the *pure palindromic automorphism group of* F_n, surjects onto $\Gamma_n[2]$, the principal level 2 congruence subgroup of $\mathrm{GL}(n, \mathbb{Z})$, via the restriction of the canonical map $\Psi : \mathrm{Aut}(F_n) \to \mathrm{GL}(n, \mathbb{Z})$. Glover-Jensen [31] attribute this surjection to Collins [18], although it is not made explicit in Collins' paper that the restriction of Ψ is onto. We show that this is indeed the case in Chapter 2, and obtain the short exact sequence

$$1 \longrightarrow \mathcal{PI}_n \longrightarrow \mathrm{P}\Pi\mathrm{A}_n \longrightarrow \Gamma_n[2] \longrightarrow 1,$$

where $\mathcal{PI}_n$ is the group $\mathrm{IA}_n \cap \mathrm{P}\Pi\mathrm{A}_n$, which we call the *palindromic Torelli group*.

One particularly strong motivation to study $\Pi\mathrm{A}_n$ arises from the extensive analogy between $\mathrm{Aut}(F_n)$ and the mapping class group $\mathrm{Mod}(S)$ of a closed, oriented surface S. The *hyperelliptic mapping class group* $\mathrm{SMod}(S)$ is the centraliser in $\mathrm{Mod}(S)$ of a fixed *hyperelliptic involution*, s, that is, a member of $\mathrm{Mod}(S)$ that acts as $-I$ on $H_1(S, \mathbb{Z})$. The obvious analogue of s in $\mathrm{Aut}(F_n)$ is the automorphism ι that inverts each $x \in X$; then clearly ι acts as $-I$ on $H_1(F_n, \mathbb{Z})$. The best candidate then for an analogy of $\mathrm{SMod}(S)$ in $\mathrm{Aut}(F_n)$ is the centraliser of ι: this is precisely $\Pi\mathrm{A}_n$ [31]. Thus, by studying $\Pi\mathrm{A}_n$ we may extend the analogy that holds between $\mathrm{Aut}(F_n)$ and $\mathrm{Mod}(S)$. We explore this analogy in further detail in Chapter 2.

A striking comparison may be drawn between $\Pi\mathrm{A}_n$ and the *pure symmetric automorphism group of* F_n, $\mathrm{P}\Sigma A_n$, which consists of automorphisms of F_n that take each $x \in X$ to a conjugate of itself. As Collins pointed out [18], there is a finite index torison-free subgroup of $\Pi\mathrm{A}_n$, $\mathrm{E}\Pi\mathrm{A}_n$, which has a finite presentation (given in Chapter 2) extremely similar to that of $\mathrm{P}\Sigma A_n$. This similarity is not entirely surprising, as in some sense we may think of a palindrome xyx ($x, y \in X$) as a 'mod 2' version of the conjugate xyx^{-1}. One notable difference between $\Pi\mathrm{A}_n$ and $\mathrm{P}\Sigma A_n$, however, is that $\mathrm{P}\Sigma A_n$ is a subgroup of IA_n, whereas the palindromic Torelli group $\mathcal{PI}_n$ is a proper subgroup of $\Pi\mathrm{A}_n$.

In Chapter 2, we obtain an infinite generating set for $\mathcal{PI}_n$. In particular, we show that $\mathcal{PI}_n$

is the normal closure in ΠA_n of two elements. Let $P_{ij} \in \Pi A_n$ denote the automorphism mapping x_i to $x_j x_i x_j$ and fixing the other members of X $(i \neq j)$.

Theorem. *For $n \geq 3$, the group $\mathcal{PI}_n$ is normally generated in ΠA_n by the automorphisms $[P_{12}, P_{13}]$ and $(P_{23}{P_{13}}^{-1}P_{31}P_{32}P_{12}{P_{21}}^{-1})^2$.*

As an immediate corollary of this theorem, we obtain an explicit finite presentation of $\Gamma_n[2]$, induced by Collins' finite presentation of $P\Pi A_n$. We note that a version of this presentation was obtained independently by Margalit-Putman [9, p5] and R. Kobayashi [39].

To obtain this generating set, we adapt a method of Day-Putman [24]. One key tool in the proof is a Birman exact sequence for $P\Pi A_n$, which allows us to induct on n. Let

$$P\Pi A_n(k) := \{\alpha \in P\Pi A_n \mid \alpha(x_i) = x_i \text{ for } 1 \leq i \leq k\}.$$

The Birman exact sequence we establish is the short exact sequence

$$1 \longrightarrow \mathcal{J}_n(k) \longrightarrow P\Pi A_n(k) \longrightarrow P\Pi A_{n-k} \longrightarrow 1,$$

where $\mathcal{J}_n(k)$ is the appropriately defined Birman kernel. We also require finite generating sets for the stabiliser subgroups $P\Pi A_n(k)$.

Theorem. *Fix $0 \leq k \leq n$, and let $\Pi A_n(k)$ consist of automorphisms which fix $x_1, \ldots, x_k$, with the convention that $\Pi A_n(0) = \Pi A_n$. Then $\Pi A_n(k)$ is generated by its intersection with Collins' generating set for ΠA_n.*

Note that in the case $k = 0$, our proof recovers Collins' original generating set for $P\Pi A_n$ [18]. While Collins takes a purely combinatorial approach, our proof is more geometric, using Stallings' graph folding algorithm [55] to write any $\alpha \in P\Pi A_n(k)$ as a product of simple generators. The use of Stallings' algorithm was motivated by a proof of Wade [58, Theorem 4.1], which showed that the pure symmetric automorphism group $P\Sigma A_n$ is amenable to folding.

We introduce a second key tool, the *complex of partial π-bases of F_n*, denoted $\mathfrak{B}_n^\pi$, in Section 2.3. The groups ΠA_n and $\mathcal{PI}_n$ act on $\mathfrak{B}_n^\pi$, and it is this action that allows us to determine the generating set for $\mathcal{PI}_n$. If the complexes $\mathfrak{B}_n^\pi$ and $\mathfrak{B}_n^\pi/\mathcal{PI}_n$ are sufficiently highly-connected, a construction of Armstrong [2] allows us to conclude that $\mathcal{PI}_n$ is generated by its vertex stabilisers of the action on $\mathfrak{B}_n^\pi$. We obtain the following connectivity result for $\mathfrak{B}_n^\pi$.

Theorem. *For $n \geq 3$, the complex $\mathfrak{B}_n^\pi$ is simply-connected.*

The quotient $\mathfrak{B}_n^\pi/\mathcal{PI}_n$ is related to complexes already studied by Charney [14], and from Charney's work we obtain that the quotient is sufficiently connected for us to apply Armstrong's construction when $n > 3$. For the $n = 3$ case, which forms the base case of our inductive proof, the quotient is not simply-connected, so we approach the problem differently, obtaining a compatible finite presentation of the congruence group $\Gamma_3[2]$, whose relators may be lifted to a normal generating set for $\mathcal{PI}_3$. This is done in Section 2.5.

Automorphisms of right-angled Artin groups. A *right-angled Artin group* A_Γ is a finitely presented group, which may be presented so that its only relators are commutators between members of its generating set. This commuting information may be encoded using the finite simplicial graph Γ with a vertex for each generator and an edge between two vertices whenever the corresponding generators commute. Right-angled Artin groups were first studied by Baudisch [5], under the name *semifree groups*, and for completeness we note that they are also known as *partially commutative groups*, *graph groups* and *trace groups* [26]. While they are exceptionally easy to define, right-angled Artin groups provide a rich collection of complicated objects to study. For instance, at first glance, one might guess that any subgroup of A_Γ will also be a right-angled Artin group. However, in reality we observe an incredibly diverse subgroup structure. Right-angled Artin groups contain, among others, almost all surface groups [20], graph braid groups [20] and virtual 3-manifold groups. The presence of virtual 3-manifold groups as subgroups, in particular, was a crucial piece of Agol's groundbreaking proof of the Virtual Haken and Virtual Fibering Conjectures of hyperbolic 3-manifold theory [1], [60].

A further reason right-angled Artin groups are worthy of study is that they allow us to interpolate between many classes of well-studied groups. These interpolations all stem from the fact that at one extreme, when A_Γ has no relators, it is a free group, F_n, whereas at the other, when A_Γ has *all* possible relators, it is a free abelian group, $\mathbb{Z}^n$. We are thus able to interpolate between free and free abelian groups by adding or removing relators to obtain a sequence of right-angled Artin groups. Many properties shared by free and free abelian groups are shared by all right-angled Artin groups: for example, for any graph Γ, the group A_Γ is linear [22] and biautomatic [34].

The automorphism group $\mathrm{Aut}(A_\Gamma)$ of a right-angled Artin group A_Γ is also a well-studied object, as passing to automorphism groups during the aforementioned interpolation between F_n and $\mathbb{Z}^n$ allows us to interpolate between $\mathrm{Aut}(F_n)$ and $\mathrm{Aut}(\mathbb{Z}^n) = \mathrm{GL}(n, \mathbb{Z})$. The groups $\mathrm{Aut}(F_n)$ and $\mathrm{GL}(n, \mathbb{Z})$ are fundamental objects of study in geometric group theory, with numerous strong analogies holding between the two. Unifying their study in the more general context of automorphism groups of right-angled Artin groups is thus an attractive proposition. In this direction, Laurence [41], proving a conjecture of Servatius [54], obtained a finite generating set for $\mathrm{Aut}(A_\Gamma)$, and Day [25] later found a finite presentation of $\mathrm{Aut}(A_\Gamma)$. Recently, Charney-Stambaugh-Vogtmann [16] constructed a virtual classifying space for a right-angled Artin group's *outer* automorphism group, $\mathrm{Out}(A_\Gamma)$, generalising Culler-Vogtmann's so-called *outer space* of the outer automorphism group of a free group [21]. Outer space is a contractible cell complex on which $\mathrm{Out}(F_n)$ acts cocompactly with finite stabilisers. There is an analogous *auter space*, on which the group $\mathrm{Aut}(F_n)$ acts. Both spaces are free group analogues of the *Teichmüller space* of an orientable surface, and points in the spaces correspond to homotopy equivalences between graphs with fundamental group F_n.

One property shared by both $\mathrm{Aut}(F_n)$ and $\mathrm{GL}(n, \mathbb{Z})$ is that both $\mathrm{Out}(\mathrm{Aut}(F_n))$ and $\mathrm{Out}(\mathrm{GL}(n, \mathbb{Z}))$ are finite. We interpret this as 'algebraic rigidity': up to conjugation, all but finitely many of the automorphisms of these groups are induced by the conjugation action of the group on itself. Dyer-Formanek [27] showed that $\mathrm{Out}(\mathrm{Aut}(F_n)) = 1$, as did Bridson-Vogtmann [10], using more geometric methods, as well as Khramtsov [38]. (Bridson-Vogtmann and Khramtsov also showed that $\mathrm{Out}(\mathrm{Out}(F_n)) = 1$ for $n \geq 3$). Hua-Reiner [35] explicitly computed $\mathrm{Out}(\mathrm{GL}(n, \mathbb{Z}))$, its structure depending, in general, on the parity of n. They found that for all n, the order of $\mathrm{Out}(\mathrm{GL}(n, \mathbb{Z}))$ is at most 4. We thus say that the orders of $\mathrm{Out}(\mathrm{Aut}(F_n))$, $\mathrm{Out}(\mathrm{Out}(F_n))$ and $\mathrm{Out}(\mathrm{GL}(n, \mathbb{Z}))$ are *uniformly* bounded above for all n by 4. In Chapter 3, we show that no such uniform upper bound exists when we consider a larger class of right-angled Artin groups.

Theorem. *For any $N \in \mathbb{N}$, there exists a right-angled Artin group A_Γ such that*

$$|\mathrm{Out}(\mathrm{Aut}(A_\Gamma))| > N.$$

We prove this theorem in two ways: our first proof uses right-angled Artin groups with non-trivial centre, while in our second proof, we work over right-angled Artin groups with trivial centre. We also prove the analogous result for $\mathrm{Out}(A_\Gamma)$.

Theorem. *For any $N \in \mathbb{N}$, there exists a right-angled Artin group A_Γ such that*

$$|\mathrm{Out}(\mathrm{Out}(A_\Gamma))| > N.$$

Our strategy for proving both of these theorems is to place certain constraints upon the graph Γ. The structure of $\mathrm{Aut}(A_\Gamma)$ and $\mathrm{Out}(A_\Gamma)$ heavily depends upon the structure of Γ: the constraints we place upon Γ lead to tractable decompositions of these groups as semi-direct products. We exploit these decompositions to construct many explicit examples of non-trivial members of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ and $\mathrm{Out}(\mathrm{Out}(A_\Gamma))$, proving the theorems.

These two theorems fit into a more general framework of algebraic rigidity within geometric group theory. For instance, the outer automorphism groups of many mapping class groups and braid groups is $\mathbb{Z}/2$ [28], [36]. In keeping with these results, and those of Hua-Reiner on $\mathrm{GL}(n, \mathbb{Z})$, further inspection of the members of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ we construct in Chapter 3 shows that they generate a direct sum of finitely many copies of $\mathbb{Z}/2$.

An open question is whether or not there exist infinite order members of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ and $\mathrm{Out}(\mathrm{Out}(A_\Gamma))$, as our methods only yield finite order elements. We state the following ambitious problem.

Problem. Classify the graphs Γ for which $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ (resp. $\mathrm{Out}(\mathrm{Out}(A_\Gamma))$) is (i) trivial, (ii) finite, and (iii) infinite.

1.1 Conventions

Throughout this thesis, we shall apply functions from right to left. For $g, h \in G$ a group, we let $[g, h] = ghg^{-1}h^{-1}$ be the commutator of g and h, and we write $g^h = hgh^{-1}$. When it is unambiguous, we shall conflate a relation $P = Q$ in a group with its relator PQ^{-1}.

In general, we shall think of a graph Y as a one-dimensional CW complex. Edges shall be oriented, with the reverse of an edge e being denoted $\bar{e}$, however we shall frequently forget about this orientation. Explicitly, an *orientation* of Y is a set containing exactly one of

e or $\bar{e}$ for each edge e of Y. When we refer to the *underlying unoriented graph of* Y, we mean the CW complex taken without orientations on the edges. Given an (oriented) edge e, we denote by $i(e)$ and $t(e)$ the initial and terminal vertices of e, respectively. We will frequently represent the edge e using the notation

$$i(e) - t(e).$$

A path in Y is taken to be a sequence of edges of Y

$$f_1 f_2 \dots f_k$$

such that $t(f_i) = i(f_{i+1})$, for $1 \leq i < k$. A path is said to be *reduced* if $f_i \neq \overline{f_{i+1}}$ for $1 \leq i < k$. Note that we may sensibly talk about the orientation of a path p, and define $\bar{p}$ to be the reverse of the path p. The *fundamental group of* Y *based at* b, denoted $\pi_1(Y, b)$, is defined to be the set paths beginning and ending at b, up to insertion and deletion of subpaths of the form $e\bar{e}$ (e an edge of Y), with multiplication defined by composition of paths.

A *map of (oriented) graphs* $\theta : Y \to Z$ is a map taking edges to edges and vertices to vertices that preserves the structure of Y in the obvious way. Such a map induces a homomorphism

$$\theta_* : \pi_1(Y, b) \to \pi_1(Z, \theta(b)).$$

Chapter 2

Palindromic automorphisms of free groups

2.1 Introduction

Let F_n be the free group of rank n on some fixed free basis X. A *palindrome* on X is a word on $X^{\pm 1}$ that reads the same backwards as forwards. The *palindromic automorphism group of* F_n, denoted $\Pi\mathrm{A}_n$, consists of automorphisms of F_n that take each member of X to a palindrome. Collins [18] introduced the group $\Pi\mathrm{A}_n$ in 1995 and proved that it is finitely presented, giving an explicit presentation. Glover-Jensen [31] obtained further results about $\Pi\mathrm{A}_n$, utilising a contractible subspace of the so-called 'auter space' of F_n on which $\Pi\mathrm{A}_n$ acts cocompactly and with finite stabilisers. For instance, they are able to calculate that the virtual cohomological dimension of $\Pi\mathrm{A}_n$ is $n-1$. One reason in particular that $\Pi\mathrm{A}_n$ is of interest to geometric group theorists is that it is an obvious free group analogue of the symmetric mapping class group of an oriented surface, a connection we shall further discuss later in this section.

Recall that the *Torelli group* of $\mathrm{Aut}(F_n)$, denoted IA_n, is the kernel of the canonical surjection $\mathrm{Aut}(F_n) \to \mathrm{GL}(n, \mathbb{Z})$. The group IA_n is very well-studied, however there are still many open questions regarding its structure and properties. In this chapter, we are primarily concerned with the intersection of $\Pi\mathrm{A}_n$ with IA_n. We denote this intersection by $\mathcal{PI}_n$, and refer to it as the *palindromic Torelli group of* F_n. Little appears to be known about the group $\mathcal{PI}_n$: Collins [18] first pointed that it is non-trivial, and Jensen-McCammond-Meier [37, Corollary 6.3] showed that $\mathcal{PI}_n$ is not homologically finite for $n \geq 3$. The main theorem of this chapter establishes a generating set for $\mathcal{PI}_n$. We let $P_{ij} \in \Pi\mathrm{A}_n$ denote the automorphism mapping x_i to $x_j x_i x_j$ for $x_i, x_j \in X$ $(i \neq j)$ and fixing all other members of X.

Theorem 2.1.1. *The group $\mathcal{PI}_n$ is normally generated in $\Pi\mathrm{A}_n$ by the automorphisms $[P_{12}, P_{13}]$ and $(P_{23}{P_{13}}^{-1}P_{31}P_{32}P_{12}{P_{21}}^{-1})^2$.*

Let $\Gamma_n[2]$ denote the principal level 2 congruence subgroup of $\mathrm{GL}(n, \mathbb{Z})$: that is, the kernel of the map $\mathrm{GL}(n, \mathbb{Z}) \to \mathrm{GL}(n, \mathbb{Z}/2)$ that reduces matrix entries mod 2. The palindromic Torelli group forms the kernel of a short exact sequence with quotient $\Gamma_n[2]$, discussed in Chapter 2.2. For $1 \leq i \neq j \leq n$, let $S_{ij} \in \mathrm{GL}(n, \mathbb{Z})$ have 1s on the diagonal and 2 in the (i, j) position, with 0s elsewhere, and let $O_i \in \mathrm{GL}(n, \mathbb{Z})$ differ from the identity only in having -1 in the (i, i) position. Theorem 2.1.1 has the following corollary. Note that for $n = 2$ and $n = 3$, some of these relators do not exist: in these cases, we simply remove them to obtain a complete list of defining relators.

Corollary 2.1.2. *The principal level 2 congruence group $\Gamma_n[2]$ of $\mathrm{GL}(n, \mathbb{Z})$ is generated by*

$$\{S_{ij}, O_i \mid 1 \leq i \neq j \leq n\},$$

subject to the defining relators

1. ${O_i}^2$,
2. $[O_i, O_j]$,
3. $(O_i S_{ij})^2$,
4. $(O_j S_{ij})^2$,
5. $[O_i, S_{jk}]$,
6. $[S_{ki}, S_{kj}]$,
7. $[S_{ij}, S_{kl}]$,
8. $[S_{ji}, S_{ki}]$,
9. $[S_{kj}, S_{ji}]S_{ki}^{-2}$,
10. $(S_{ij}{S_{ik}}^{-1}S_{ki}S_{ji}S_{jk}{S_{kj}}^{-1})^2$

where $1 \leq i, j, k, l \leq n$ are pairwise distinct.

We note that in the proof of Theorem 2.1.1 and Corollary 2.1.2, it becomes apparent that not every relator of type 10 is needed: in fact, for each choice of three indices i, j and k, we need only select one such word (and disregard the others, for which the indices have been permuted).

Corollary 2.1.2 gives a particularly natural presentation for $\Gamma_n[2]$ [47], as the relations which hold between the S_{ij} bear a strong resemblance to the Steinberg relations which hold between the transvections generating $\mathrm{SL}(n,\mathbb{Z})$, as we now explain. Let E_{ij} be the elementary matrix with 1 in the (i,j) position. Clearly $S_{ij} = {E_{ij}}^2$. A complete set of relators for the group $\langle E_{ij}\rangle = \mathrm{SL}(n,\mathbb{Z})$ $(n \geq 3)$ is

1. $[E_{ij}, E_{ik}]$,
2. $[E_{ik}, E_{jk}]$,
3. $[E_{ij}, E_{jk}]{E_{ik}}^{-1}$,
4. $(E_{12}{E_{21}}^{-1}E_{12})^4$,

where the indices i, j, k are taken to be pairwise distinct. Relators of type $1-3$ are referred to as *Steinberg relations* [47, §5]. As pointed out by Margalit-Putman [45], the relations holding between the S_{ij} consist of 'Steinberg-like' relations (types $6-9$ in Corollary 2.1.2) and one extra relation (relator 10), which bears a certain resemblance to the relator $(E_{12}{E_{21}}^{-1}E_{12})^4$. A similar presentation for $\Gamma_n[2]$ was obtained independently by Kobayashi [39], and was also known to Margalit-Putman [45].

2.1.1 A comparison with mapping class groups

While $\Pi\mathrm{A}_n$ is defined entirely algebraically, it may viewed as a free group analogue of a group that arises in low-dimensional topology. Let S_g^1 be the compact, connected, oriented surface of genus g with one boundary component. Recall that the *mapping class group of* S_g^1, denoted $\mathrm{Mod}(S_g^1)$, is the group of orientation-preserving homeomorphisms up to isotopy. Our convention is only to consider homeomorphisms and isotopies that fix the boundary component point-wise. The mapping class group has induced actions on both the fundamental group $\pi_1(S_g^1) = F_{2g}$ and the first homology group $H_1(S_g^1,\mathbb{Z}) = \mathbb{Z}^{2g}$ of the surface. Both of these actions shall be of interest to us.

Let S_g be the result of capping off the boundary component of S_g^1 with a disk. A *hyperelliptic*

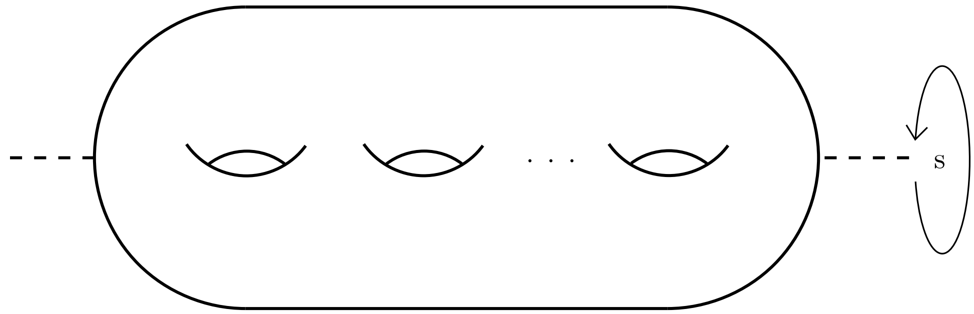

Figure 2.1: *The hyperelliptic involution $s \in \mathrm{Mod}(S_g)$ shown rotates the surface by π radians along the indicated axis.*

involution of S_g is an involution $s \in \mathrm{Mod}(S_g)$ that acts as $-I$ on $H_1(S_g, \mathbb{Z})$. For $g \geq 1$, all hyperelliptic involutions are conjugate in $\mathrm{Mod}(S_g)$ [29, Proposition 7.15]: an example of one is seen in Figure 2.1. As the disk we attached to obtain S_g is invariant under this involution s, we may also consider the involution s shown in Figure 2.1 as a homeomorphism of S_g^1, however notice that it does not fix the boundary component point-wise. Clearly, we still have $s \in \mathrm{Homeo}^+(S_g^1)$, the group of orientation-preserving self-homeomorphisms of S_g^1.

We define the *hyperelliptic mapping class group of* S_g^1, denoted $\mathrm{SMod}(S_g^1)$, to be the subgroup of $\mathrm{Mod}(S_g^1)$ of mapping classes that have a representative that commute with s in $\mathrm{Homeo}^+(S_g^1)$. There is an analogously-defined *hyperelliptic mapping class group of* S_g, denoted $\mathrm{SMod}(S_g)$, with a more succinct definition: it is simply the centraliser of $[s]$ in $\mathrm{Mod}(S_g)$, where $[s]$ is the isotopy class of the involution $s \in \mathrm{Homeo}^+(S_g)$. Recall that, like $\mathrm{Aut}(F_n)$, $\mathrm{Mod}(S_g)$ and $\mathrm{Mod}(S_g^1)$ have large subgroups that act trivially on first homology of the surface. These groups are also called *Torelli groups*, and are denoted $\mathcal{I}_g$ and $\mathcal{I}_g^1$, respectively.

Translating these notions into the context of $\mathrm{Aut}(F_n)$, an obvious analogue in $\mathrm{Aut}(F_n)$ of the involution s is the automorphism ι that inverts each member of the free basis X. The following proposition, which is noted by Glover-Jensen [31], establishes that $\Pi\mathrm{A}_n$ is the centraliser of ι in $\mathrm{Aut}(F_n)$.

Proposition 2.1.3. *The centraliser in* $\mathrm{Aut}(F_n)$ *of* ι *is* $\Pi\mathrm{A}_n$.

Proof. We carry out a straightforward calculation. Let $\alpha \in \mathrm{Aut}(F_n)$, $x \in X$ and write $\alpha(x) = w_1 \dots w_r$ (for some $r \in \mathbb{N}$ and $w_i \in X^{\pm 1}$). The automorphism α centralises ι if and

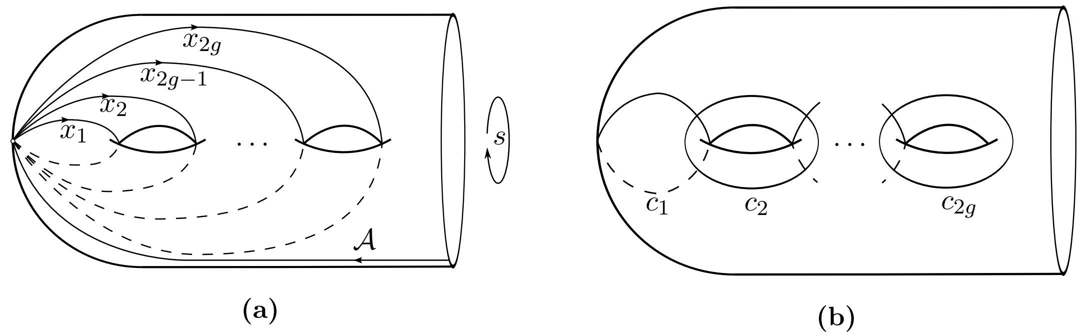

Figure 2.2:

(a) *The involution s rotates the surface by π radians. Under the classical Nielsen embedding, we may view the braid group $B_{2g} \leq \mathrm{SMod}(S_g^1)$ as a subgroup of $\Pi\mathrm{A}_{2g} \leq \mathrm{Aut}(F_{2g})$, where F_{2g} is the free group on the oriented loops $x_1, \dots, x_{2g}$.*

(b) *The standard symmetric chain in S_g^1. The Dehn twists about $c_1, \dots, c_{2g}$ generate $\mathrm{SMod}(S_g^1) \cong B_{2g+1}$.*

only $\alpha\iota = \iota\alpha$: that is, if and only if

$$w_r^{-1} \dots w_1^{-1} = w_1^{-1} \dots w_r^{-1}.$$

Assuming, without loss of generality, that $w_1 \dots w_r$ was a reduced expression of $\alpha(x)$, we have that $\alpha(x)$ is a palindrome, and so the proposition is established. $\square$

The comparison between $\Pi\mathrm{A}_n$ and $\mathrm{SMod}(S_g^1)$ is made more precise using the classical Nielsen embedding $\mathrm{Mod}(S_g^1) \hookrightarrow \mathrm{Aut}(F_{2g})$. Take the $2g$ oriented loops seen in Figure 2.2a as a free basis for $\pi_1(S_g^1)$. Observe that s acts on these loops by switching their orientations. In order to use Nielsen's embedding into $\mathrm{Aut}(F_{2g})$, we must take these loops to be based on the boundary; we surger using the arc $\mathcal{A}$ to achieve this. The group $\mathrm{SMod}(S_g^1)$ is isomorphic to the braid group B_{2g+1} by the Birman-Hilden theorem [8], and is generated by Dehn twists about the curves in the standard, symmetric chain on S_g^1, seen in Figure 2.2b. The Dehn twists about the $2g-1$ curves $c_2, \dots, c_{2g}$ generate the braid group B_{2g}. Taking the loops seen in Figure 2.2a as our free basis X, a straightforward calculation shows that the images of these $2g-1$ twists in $\mathrm{Aut}(F_{2g})$ lie in $\Pi\mathrm{A}_{2g}$. Specifically, the twist about c_{i+1}

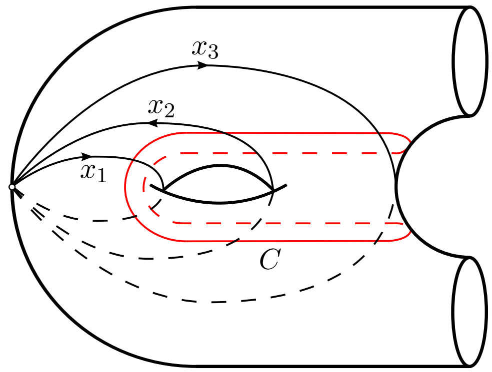

Figure 2.3: *The Dehn twist about the symmetric, separating curve C is the preimage in $\mathcal{SI}(S_g^1)$ of $\chi \in \mathcal{PI}_{2g}$ under the Nielsen embedding.*

is taken to the automorphism Q_i of the form

$$\begin{aligned} x_i &\mapsto x_{i+1}, \\ x_{i+1} &\mapsto x_{i+1}{x_i}^{-1}x_{i+1}, \\ x_j &\mapsto x_j \end{aligned}$$

for $1 \leq i < 2g$ and $j \neq i, i+1$. This shows that $\Pi\mathrm{A}_n$ contains the braid group B_n as a subgroup, when n is even. This embedding of B_n is a restriction of one studied by Perron-Vannier [51] and Crisp-Paris [19]. When n is odd, we also have $B_n \hookrightarrow \Pi\mathrm{A}_n$, since discarding Q_1 gives a generating set for B_{2g-1} inside $\Pi\mathrm{A}_{2g-1} \leq \mathrm{Aut}(F_{2g})$.

The main focus of our study of this chapter is the palindromic Torelli group, $\mathcal{PI}_n$. This group arises as a natural analogue of a subgroup of $\mathrm{SMod}(S_g^1)$. The *Torelli subgroup* of $\mathrm{Mod}(S_g^1)$, denoted $\mathcal{I}_g^1$, consists of mapping classes that act trivially on $H_1(S_g^1, \mathbb{Z})$. There is non-trivial intersection between $\mathcal{I}_g^1$ and $\mathrm{SMod}(S_g^1)$; we define $\mathcal{SI}(S_g^1) := \mathrm{SMod}(S_g^1) \cap \mathcal{I}_g^1$ to be the *hyperelliptic Torelli group*. Brendle-Margalit-Putman [9] recently proved a conjecture of Hain [32], also stated by Morifuji [48], showing that $\mathcal{SI}(S_g^1)$ is generated by Dehn twists about separating simple closed curves of genus 1 and 2 that are fixed by s. (Recall that a simple closed curve c on a surface S is said to be *separating* if $S \setminus c$ is disconnected, and that the *genus* of such a curve c is the minimum of the genera of the connected components of $S \setminus c$). Our generating set for $\mathcal{PI}_n$ compares favourably with Brendle-Margalit-Putman's for $\mathcal{SI}(S_g^1)$, in the following way. The generator $\chi := (P_{23}{P_{13}}^{-1}P_{31}P_{32}P_{12}{P_{21}}^{-1})^2$ in the statement of Theorem 2.1.1 can be realised topologically on S_g^1, as it lies in the image of $\mathcal{SI}(S_g^1)$ in $\Pi\mathrm{A}_{2g}$. Direct computation shows that χ is the image of the Dehn twist about the

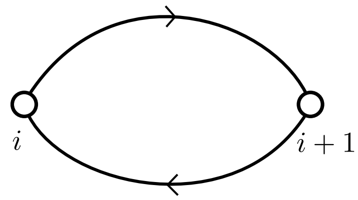

Figure 2.4: *The standard braid generator σ_i $(1 \leq i < 2g+1)$ interchanges the ith and $(i+1)$th punctures in a clockwise direction, as shown.*

curve C seen in Figure 2.3, with the loops oriented as shown. Note that C is a symmetric, separating curve of genus 1, and so is one of the two normal generators of Brendle-Margalit-Putman's generating set. We shall see in Proposition 2.3.7 that conjugates in $\Pi\mathrm{A}_n$ of our other normal generator $[P_{12}, P_{13}]$ do not suffice to generate $\mathcal{PI}_n$, so we observe that our generating set involves Brendle-Margalit-Putman's generators in a significant way. The similarity between $\mathcal{SI}_g^1$ and $\mathcal{PI}_n$ is not just a superficial comparison of definitions: the Nielsen embedding gives rise to a deeper connection between these two groups.

The analogy breaks down. One way in which the analogy between $\mathcal{PI}_n$ and $\mathcal{SI}(S_g^1)$ breaks down, however, is their behaviour when $\Pi\mathrm{A}_n$ and $\mathrm{SMod}(S_g^1)$ are abelianised, to $\mathbb{Z}/2$ and $\mathbb{Z}$ respectively. An immediate corollary of Theorem 2.1.1 is that $\mathcal{PI}_n$ vanishes in the abelianisation of $\Pi\mathrm{A}_n$. In contrast, the image of $\mathcal{SI}(S_g^1)$ in the abelianisation of $\mathrm{SMod}(S_g^1)$ is $4\mathbb{Z}$, which we now prove.

Theorem 2.1.4. *The group $\mathcal{SI}(S_g^1)$ has image $4\mathbb{Z}$ in the abelianisation of* $\mathrm{SMod}(S_g^1)$*.*

Proof. We pass to the $(2g+1)$-punctured disk of which S_g^1 is a branched double cover by the involution s, and use the Birman-Hilden theorem to identify $\mathrm{SMod}(S_g^1)$ with the braid group B_{2g+1}. We refer the reader to Farb-Margalit [29, Chapter 9.4] for a detailed discussion of this procedure.

Let σ_i denote the standard half-twist generator of B_{2g+1} that swaps the ith and $(i+1)$th punctures in a clockwise direction, as seen in Figure 2.4. A Dehn twist about a genus 1 (resp. 2) symmetric separating curve in S_g^1 descends to the square of a Dehn twist about a simple closed curve in D_{2g+1} surrounding 3 (resp. 5) punctures. A straightforward calculation shows that

$$T_3 := \sigma_1^2[\sigma_2\sigma_1^2\sigma_2],$$

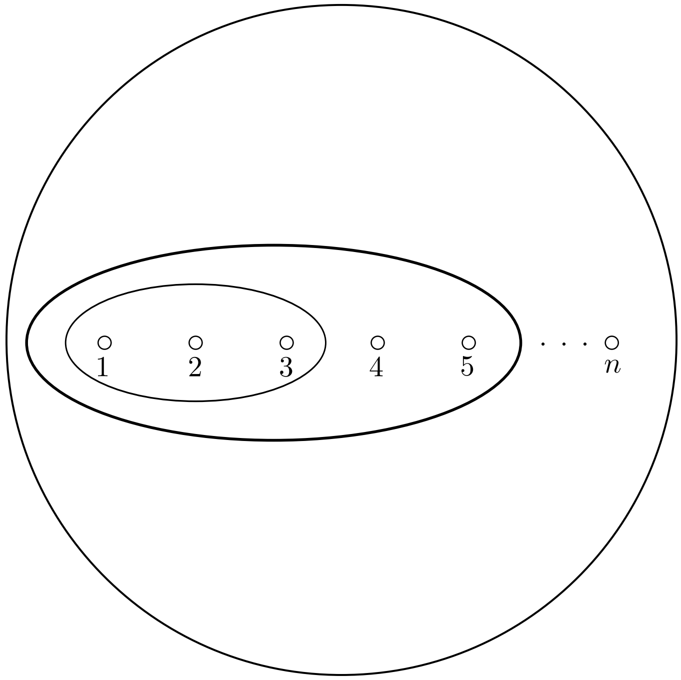

Figure 2.5: *Curves in a punctured disk surrounding 3 and 5 punctures, respectively. For $n = 2g+1$, Brendle-Margalit-Putman show that the squares of the Dehn twists about these curves normally generate the image of $\mathcal{SI}_{2g+1}$ in B_{2g+1}.*

and

$$T_5 := \sigma_1^2[\sigma_2\sigma_1^2\sigma_2][\sigma_3\sigma_2\sigma_1^2\sigma_2\sigma_3][\sigma_4\sigma_3\sigma_2\sigma_1^2\sigma_2\sigma_3\sigma_4],$$

are equal to Dehn twists about the simple closed curves in D_{2g+1} surrounding 3 and 5 punctures, respectively, shown in Figure 2.5. The image of $\mathcal{SI}(S_g^1)$ in the abelianisation of B_{2g+1} depends only upon the images of T_3 and T_5, as their squares normally generate $\mathcal{SI}(S_g^1)$.

Let $\mathbb{Z} = \langle t \rangle$ be the abelianisation of B_{2g+1}. The image in $\mathbb{Z}$ of ${T_3}^2$ is t^{12}, and the image of ${T_5}^2$ is t^{40}, so $\mathcal{SI}(S_g^1)$ has image $\langle t^4 \rangle = 4\mathbb{Z}$. □

We also observe that Dehn twists about both genus 1 and genus 2 separating curves are needed to generate $\mathcal{SI}(S_g^1)$, as we show in the following corollary.

Corollary 2.1.5. *The set of Dehn twists about symmetric simple separating curves of genus 1 (resp. 2) does not generate $\mathcal{SI}(S_g^1)$.*

Proof. The subgroup normally generated by only twists about genus 1 (resp. 2) curves has image $12\mathbb{Z}$ (resp. $40\mathbb{Z}$) in the abelianisation of B_{2g+1}, and so cannot equal $\mathcal{SI}(S_g^1)$. $\square$

2.1.2 Approach of the proof of Theorem 2.1.1

To prove Theorem 2.1.1, we employ a standard technique of geometric group theory: we find a sufficiently connected simplicial complex on which $\mathcal{PI}_n$ acts with sufficiently connected quotient, and use a theorem of Armstrong [2] to conclude that $\mathcal{PI}_n$ ($n > 3$) is generated by the action's vertex stabilisers. This approach is modelled on a proof of Day-Putman [24] which recovers Magnus' finite generating set for the Torelli subgroup of $\mathrm{Aut}(F_n)$. We treat the $n = 3$ case separately, obtaining a compatible finite presentation for $\Gamma_3[2]$, whose relators correspond to a normal generating set for $\mathcal{PI}_3$ in $\Pi\mathrm{A}_3$.

2.1.3 Outline of chapter

In Section 2.2, the definitions of the palindromic automorphism group and palindromic Torelli group of a free group are given, along with some elementary properties of these groups. In Section 2.3, we introduce our new complex, the complex of partial π-bases of F_n, and use it to obtain a generating set for $\mathcal{PI}_n$. In Section 2.4, we prove key results about the connectivity of the complexes involved in the proof of Theorem 2.1.1. In Section 2.5, we obtain a finite presentation of $\Gamma_3[2]$ used in the base case of our inductive proof of Theorem 2.1.1.

2.2 The palindromic automorphism group

Let F_n be the free group of rank n, on some fixed free basis $X := \{x_1, \ldots, x_n\}$.

2.2.1 Palindromes in F_n

For a word $w = l_1 \ldots l_k$ on $X^{\pm 1}$, let w^{rev} denote the *reverse* of w; that is, we have $w^{\mathrm{rev}} = l_k \ldots l_1$. Such a word w is said to be a *palindrome* on X if $w^{\mathrm{rev}} = w$. For example, x_1, ${x_2}^2$ and $x_2 x_1^{-1} x_2$ are all palindromes on X.

An odd-length palindrome $w^{\text{rev}}x_i^{\epsilon_i}w$ ($\epsilon_i \in \{\pm 1\}$) and the conjugate $w^{-1}x_i^{\epsilon_i}w$ have the same image in the free Coxeter group quotient of F_n obtained by adding the relators ${x_i}^2 = 1$ ($1 \leq i \leq n$). We might therefore expect there to be some connection between conjugation and palindromes in F_n, however the following proposition shows that they are rather orthogonal concepts.

Proposition 2.2.1. *Let $p \in F_n$ be a palindrome.*

1. *If p has odd length, it is the only palindrome in its conjugacy class,*
2. *If p has even length, there is precisely one other palindrome $p' \neq p$ in its conjugacy class.*

Proof. Without loss of generality, we assume that p is a reduced word in F_n. This proof may seem heavy-handed, but it yields more information about palindromic conjugates of p than more elementary proofs might. We deal with the odd length case first.

Suppose that $q \in F_n$ is a palindrome conjugate to p, which is also reduced as a word in F_n. This means q is simply a cyclic permutation of the word p. Suppose p has length $2k+1$ ($k \geq 0$), and let

$$p = l_{-k} \dots l_{-1} l_0 l_1 \dots l_k,$$

where $l_i \in X^{\pm 1}$ and $l_{-i} = l_i$. We have a similar expression for q, with

$$q = \tilde{l}_{-k} \dots \tilde{l}_{-1} \tilde{l}_0 \tilde{l}_1 \dots \tilde{l}_k,$$

where $\tilde{l}_i \in X^{\pm 1}$ and $\tilde{l}_{-i} = \tilde{l}_i$. Our strategy is to find a way of translating the condition $\tilde{l}_{-i} = \tilde{l}_i$ into one between members of $\{l_i\}$.

To do this translating, we work in the ring $\mathbb{Z}/(2k+1)$, setting up the obvious bijection between the set of letters $\{l_i\}$ of p and

$$\mathbb{Z}/(2k+1) = \{-k, \dots, -1, 0, 1, \dots k\}.$$

Fix $c \in \mathbb{Z}/(2k+1)$, and suppose that $\tilde{l}_0 = l_c$. We refer the reader to the graph $\mathcal{K}$ in Figure 2.6, where vertices correspond to members of $\mathbb{Z}/(2k+1)$, and two vertices i and j are joined by an edge if $l_i = l_j$. The horizontal edges arise due to the relations $l_{-i} = l_i$, and the non-horizontal, dashed edges arise due to the relations $\tilde{l}_{-i} = \tilde{l}_i$. We obtain a closed path in

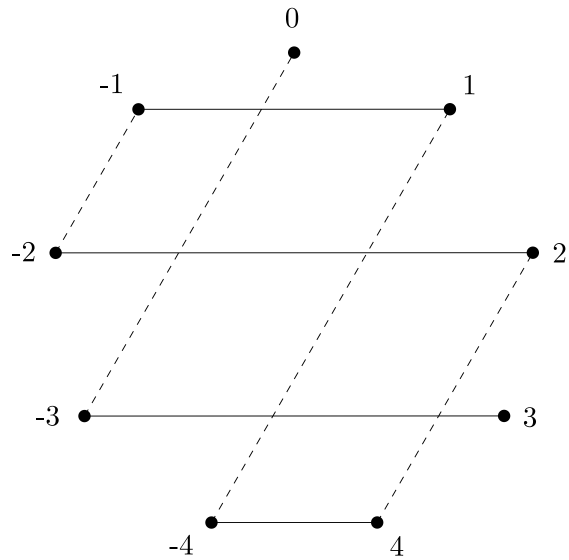

Figure 2.6: *The graph $\mathcal{K}$ for $k = 4$ and $c = 3$.*

$\mathcal{K}$ by following an alternating sequence of horizontal and dashed edges: the one exception to this is the path joining 0 and c. Clearly $l_0 = l_c$, so we add an edge between the vertices 0 and c. To traverse a horizontal edge at the vertex i, we move to the vertex $-i$: we call such a move the *negation* of a vertex. To traverse a dashed edge at the vertex i, we move to the vertex $-i + 2c$: this corresponds to 'conjugating' the negation of a vertex by the rotation $j \mapsto j + c$.

By repeatedly applying these two operations, one after the other, we see that each closed path consists precisely of the members of the cosets $i + \langle 2c \rangle$ and $-i + \langle 2c \rangle$, for some i, where $\langle 2c \rangle$ is the ideal generated by $2c$ in $\mathbb{Z}/(2k+1)$. Let d be such that $(\mathbb{Z}/(2k+1))/\langle 2c \rangle \cong \mathbb{Z}/d$. Obviously d is an (odd) divisor of $2k+1$, and p is wholly determined by $l_0, l_1, \dots, l_{d-1}$, since, up to a cyclic reordering, it is simply some power of $l_0 l_1 \dots l_{d-1}$. Since $\gcd(2c, 2k+1) \leq c$, it must be the case that c is a multiple of d. The vertex associated to $\tilde{l}_i$ is $i + c \bmod (2k+1)$, so $l_i = \tilde{l}_i$, since their associated vertices lie in the same coset of $\langle d \rangle$ in $\mathbb{Z}/(2k+1)$. Thus $p = q$.

When p is an even length palindrome of length, say, $2k$, the above argument is not applicable immediately, as there is no way to label the $2k$ vertices of the corresponding graph $\mathcal{K}$ so that traversing horizontal edges corresponds to negation in $\mathbb{Z}/2k$. We get around this by introducing $2k$ 'dummy' vertices, as seen in Figure 2.7. The labelling seen in Figure 2.7 then allows the previous argument to go through, essentially unchanged, since no dummy

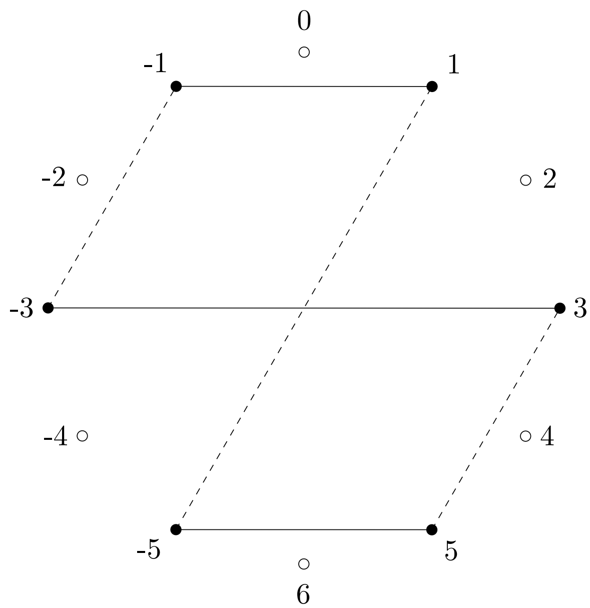

Figure 2.7: *The graph $\mathcal{K}$ for $k = 3$ and $c = 2$, where the white vertices are the dummy vertices we have introduced.*

vertex will be joined by an edge to a non-dummy vertex. Note, however, that in the even case, the palindrome $l_1 \dots l_k l_k \dots l_1$ is conjugate to the palindrome $l_k \dots l_1 l_1 \dots l_k$, which may be a different word in F_n. This corresponds to doing a 'half-rotation' of the graph $\mathcal{K}$ that is not possible in the odd case. Any other rotation leads to the same analysis as in the odd case. $\square$

2.2.2 Palindromic automorphisms of F_n

We fix the free basis $X = \{x_1, \dots, x_n\}$ once and for all. An automorphism $\alpha \in \mathrm{Aut}(F_n)$ is said to be *palindromic* if for each $x_i \in X$, the word $\alpha(x_i)$ may be written as a palindrome on X. Such automorphisms form a subgroup of $\mathrm{Aut}(F_n)$ which we call the *palindromic automorphism group of* F_n and denote by $\Pi\mathrm{A}_n$. That $\Pi\mathrm{A}_n$ is a group is easily shown by verifying that $\Pi\mathrm{A}_n$ is the centraliser in $\mathrm{Aut}(F_n)$ of the automorphism ι which inverts each member of X, as we did in the proof of Proposition 2.1.3. The following proposition allows us to conclude that the palindromes $\alpha(x_i)$ must all have odd length and each have a unique 'central' letter.

Proposition 2.2.2. *Let $\alpha \in \Pi\mathrm{A}_n$ and $x_i \in X$. Then $\alpha(x_i) = w^{\mathrm{rev}} \sigma(x_i)^{\epsilon_i} w$, where w is a word on $X^{\pm 1}$, σ is a permutation of X and $\epsilon_i \in \{\pm 1\}$.*

Proof. For a palindrome $p = w^{\text{rev}} x_i^{\epsilon_i} w \in F_n$ of odd length ($w \in F_n$, $x_i \in X$, $\epsilon_i \in \{\pm 1\}$), let $c(p) = x_i$. We refer to $c(p)$ as the *core* of p. The following argument is implicit in the work of Collins [18].

Let $\alpha \in \Pi A_n$. There is a natural surjection $F_n \to (\mathbb{Z}/2)^n$ induced by adding the relators ${x_i}^2$ and $[x_i, x_j]$ to F_n ($1 \leq i \neq j \leq n$): since $\alpha(X)$ is a free basis for F_n, its image under this surjection must suffice to generate $(\mathbb{Z}/2)^n$. If some $\alpha(x_i)$ was of even length, it would have zero image in $(\mathbb{Z}/2)^n$, and so the image of $\alpha(X)$ could not generate. Similarly, if $c(\alpha(x_i)) = c(\alpha(x_j))$ for some $i \neq j$, then $\alpha(x_i)$ and $\alpha(x_j)$ would have the same image in $(\mathbb{Z}/2)^n$, and so again $\alpha(X)$ could not generate. □

2.2.3 Stallings' graph folding algorithm

We momentarily divert our attention to a graph theoretic technique that we shall use in Section 2.2.4. Given certain fixed choices, there is a canonical way to realise any automorphism $\alpha \in \text{Aut}(F_n)$ as a map of graphs, which we describe shortly. Stallings [55] developed a powerful technique of 'folding' graphs, one application of which is to take this map of graphs and use it to factor α as a product of simpler automorphisms. This provides a geometric proof of the finite generation of $\text{Aut}(F_n)$; we shall use similar ideas to find finite generating sets for ΠA_n and certain stabiliser subgroups, in Section 2.2.4.

We remark that while we use Stallings' combinatorial description of graphs (following Serre [53]) and foldings, it is possible to view folding more topologically, regarding graphs as topological spaces and foldings as continuous maps onto quotient spaces (see Bestvina-Handel [6]), for example).

Let Y be a finite graph with a distinguished vertex b, which will act as a base point. Select a maximal tree T in Y. We orient an edge e of T by defining the initial vertex $i(e)$ to be the endpoint of e that is closer to b under the edge metric on T: denote this orientation by $\mathcal{O}(T, b)$. Choose an orientation of the edges $Y \setminus T =: \{f_1, \ldots, f_n\}$. For any vertex v in Y, we define p_v to be the unique reduced (oriented) path in T from b to v. Let

$$y_i = p_{i(f_i)} f_i \overline{p_{t(f_i)}}$$

for $1 \leq i \leq n$. The following classical theorem gives a free basis for $\pi_1(Y, b)$, given T and the chosen associated orientations.

Theorem 2.2.3 (Lyndon-Schupp [42]). *The set* $\{y_1, \ldots, y_n\}$ *is a free basis for* $\pi_1(Y, b)$. *Moreover, a sequence of edges forming a member of* $\pi_1(Y, b)$ *may be expressed in terms of this free basis by deleting any edges of* T *and replacing each* f_i *with* y_i *and each* $\overline{f_i}$ *with* ${y_i}^{-1}$.

Let $\theta : Y \to Z$ be a map of graphs. We call θ an *immersion* if for each vertex v of Y, the restriction of θ to the edges with initial vertex v is injective, and a *homotopy equivalence* if the induced homomorphism θ_* is an isomorphism of fundamental groups.

If such a map θ is not an immersion, there must exist a vertex v of Y with two edges coming out of it that have the same image in Z. The map θ hence factors through the quotient graph Y' obtained by identifying these edges (and their terminal vertices). We get induced maps $\phi : Y \to Y'$ and $\theta' : Y' \to Z$ such that $\theta = \theta'\phi$. We call this procedure *folding*, with ϕ being called the *folding map*. In an obvious way, we think of the map θ' as being closer to being an immersion than θ, as we have removed one instance of θ failing to be an immersion. The following theorem is the key ingredient to Stallings' folding algorithm.

Theorem 2.2.4 (Stallings [55]). *Suppose* X *is a finite, connected graph. Let* $\theta : Y \to Z$ *be a map of graphs. Then*

1. *If* θ *is an immersion, then*

$$\theta_* : \pi_1(Y, b) \to \pi_1(Z, \theta(b))$$

   is an injection;

2. *If* θ *is not an immersion, there is a sequence of foldings*

$$Y \xrightarrow{\phi_1} Y_1 \xrightarrow{\phi_2} \ldots \xrightarrow{\phi_k} Y_k$$

   and an immersion $\theta' : Y_k \to Z$ *such that* $\theta = \theta'\phi_k \ldots \phi_1$.

We are interested in the case where θ is a homotopy equivalence: in this case, there are only two types of folding, as seen in Figure 2.8. Let R_n denote the graph obtained by gluing together n copies of S^1 together at a base point labelled o, and let $\theta : Y \to R_n$ be a homotopy equivalence. Following Wade [58], we refer to θ along with the choices we made

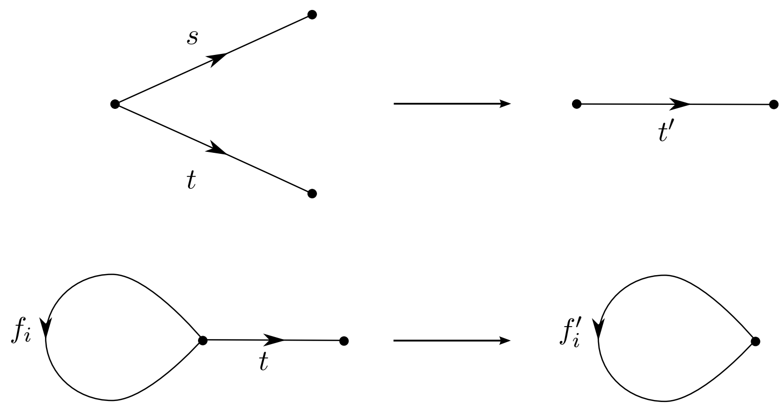

Figure 2.8: *The two types of folding that occur when θ is a homotopy equivalence. Wade [58] refers to the top fold as a type 1 fold, and to the bottom as a type 2 fold. The edges are labelled suggestively: we will demand that $s, t \in T$ and $f_i \notin T$.*

in order to state Theorem 2.2.3 (b, T, and an ordered orientation of $Y \setminus T$) as a *branding* $\mathcal{G}$ of the graph Y. With this data, Y becomes an *branded graph*, with branding $\mathcal{G}$.

Each branded graph yields an automorphism $B_{\mathcal{G}} \in \mathrm{Aut}(F_n)$. For each x_i in the free basis X of F_n, we have

$$B_{\mathcal{G}}(x_i) = \theta_*(y_i),$$

where y_i is as stated in Theorem 2.2.3, and we have made an identification between X and the (oriented, ordered) loops of R_n. Note that this is a well-defined automorphism, as we have insisted that θ is a homotopy equivalence, and so θ_* is an isomorphism.

Given a branding of Y, we may fold θ if it is not an immersion. Repeatedly folding, by Theorem 2.2.4 we eventually obtain an immersion $\theta' : Y_k \to R_n$. By observing what effect the folds of type 1 and 2 have on $B_{\mathcal{G}}$, we shall be able to write $B_{\mathcal{G}}$ as a product of what are known as *Whitehead automorphisms*, whose definition we now recall.

A *Whitehead automorphism of type 1* is simply a member of $\Omega^{\pm 1}(X)$, the group of permutations and inversions of members of X. Let $a \in X^{\pm 1}$ and $A \subset X^{\pm 1}$ be such that $a \in A$

but $a^{-1} \notin A$. The *Whitehead automorphism of type 2*, $(A, a) \in \mathrm{Aut}(F_n)$, is defined by

$$(A,a)(x_i) = \begin{cases} x_i & \text{if } x_i = a^{\pm 1} \\ ax_i & \text{if } x_i \in A \text{ and } {x_i}^{-1} \notin A \\ x_i a^{-1} & \text{if } x_i \notin A \text{ and } {x_i}^{-1} \in A \\ ax_i a^{-1} & \text{if } x_i \in A \text{ and } {x_i}^{-1} \in A \end{cases}.$$

If we insist that the edges s and t seen in Figure 2.8 lie in T, and that the edge f_i does not, carrying out either fold induces a branding $\mathcal{G}'$ of the folded graph Y' (it is non-trivial to verify that the image of T in Y' is a maximal tree; we leave this to Wade). We find that $B_{\mathcal{G}} = B_{\mathcal{G}'} \cdot W$, where W is a Whitehead automorphism of type 2. It may also be the case that we wish to carry out a fold of type 1 or type 2, but that s or t does not lie in T. Before folding, we must change maximal tree so that the relevant edges lie in the new tree. This defines a new branding $\mathcal{G}_t$ of Y. Again, we find that $B_{\mathcal{G}} = B_{\mathcal{G}_t} \cdot W$, where W is a Whitehead automorphism of type 2. With this notation set, the following propositions make these notions precise.

Proposition 2.2.5 (Proposition 3.1, [58])**.** *Suppose that we carry out a fold of type 1 to the branded graph Y, with $s, t \in T$. Then $B_{\mathcal{G}} = B_{\mathcal{G}'}$.*

To carry out a type 2 fold (that is, identify the edges t and f_i seen in Figure 2.8), first let $\epsilon = 1$ if $t \in \mathcal{O}(T, b)$ and $\epsilon = -1$ otherwise, where $\mathcal{O}(T, b)$ is the canonical orientation we assign to T.

Proposition 2.2.6 (Proposition 3.2, [58])**.** *Suppose that we carry out a fold of type 2 to the branded graph Y, with $t \in T$. Let $A \subset X^{\pm 1}$ be such that*

1. ${x_i}^{\epsilon} \in A$,
2. ${x_i}^{-\epsilon} \notin A$,
3. $x_j \in A$ *if and only if* t *or* $\bar{t}$ *is an edge of* $p_{i(f_j)}$, *and*
4. ${x_j}^{-1} \in A$ *if and only if* t *or* $\bar{t}$ *is an edge of* $p_{t(f_j)}$.

Then $B_{\mathcal{G}} = B_{\mathcal{G}'} \cdot (A, {x_i}^{\epsilon})$.

Finally, we consider the effect of changing the maximal tree T. We must do this if s or t is not in T. Without loss of generality, assume $t \notin T$. Then $t = f_j$ or $\bar{t} = f_j$, for some $1 \leq j \leq n$. Choose an edge f_j' that is contained in only one of $p_{i(f_j)}$ and $p_{t(f_j)}$ (such an edge much exist, as t has distinct endpoints. Removing f_j' from T and replacing it with t gives a new branding $\mathcal{G}_t$ of Y (again, Wade verifies that this process yields a new maximal tree). Define $\epsilon = 1$ if $f_j' \in p_{i(f_j)}$ and $\epsilon = -1$ if $f_j' \in \overline{p_{t(e_j)}}$.

Proposition 2.2.7 (Proposition 3.3, [58]). *Let $\mathcal{G}$ and $\mathcal{G}_t$ be brandings of Y as above. Let $A \subset X^{\pm 1}$ be such that*

1. ${x_j}^{\epsilon} \in A$,
2. ${x_j}^{-\epsilon} \notin A$,
3. $x_k \in A$ *if and only if* f_j' *or* $\overline{f_j'}$ *is an edge of* $p_{i(f_k)}$, *and*
4. ${x_k}^{-1} \in A$ *if and only if* f_j' *or* $\overline{f_j'}$ *is an edge of* $p_{t(f_k)}$.

Then $B_{\mathcal{G}} = B_{\mathcal{G}'} \cdot (A, {x_j}^{\epsilon})$.

By Theorem 2.2.4, we know that after a finite sequence of foldings, we obtain an immersion $\theta' : Y_k \to R_n$ that is also a homotopy equivalence. Let $\mathcal{G}'$ be any branding of the graph Y_k under θ'. Lemma 2.7 of Wade [58] allows us to conclude that θ' is a graph isomorphism and that $B_{\mathcal{G}'}$ is a Whitehead automorphism of type 1. Thus, our sequence of foldings terminates at $\theta' : Y_k \to R_n$, and we have a factorisation of $B_{\mathcal{G}}$ into Whitehead automorphisms.

2.2.4 Finite generation of $\Pi\mathrm{A}_n$

Collins first studied the group $\Pi\mathrm{A}_n$, giving a finite presentation for it. For $i \neq j$, let $P_{ij} \in \Pi\mathrm{A}_n$ map x_i to $x_j x_i x_j$ and fix x_k with $k \neq i$. For each $1 \leq j \leq n$, let $\iota_j \in \Pi\mathrm{A}_n$ map x_j to x_j^{-1} and fix x_k with $k \neq j$. We refer to P_{ij} as an *elementary palindromic automorphism* and to ι_j as an *inversion*. We let $\Omega^{\pm 1}(X)$ denote the group generated by the inversions and the permutations of X. The group generated by all elementary palindromic automorphisms and inversions is called the *pure palindromic automorphism group* of F_n, and is denoted $\mathrm{P}\Pi\mathrm{A}_n$.

Collins showed that $\Pi\mathrm{A}_n \cong \mathrm{E}\Pi\mathrm{A}_n \rtimes \Omega^{\pm 1}(X)$, where $\mathrm{E}\Pi\mathrm{A}_n = \langle P_{ij} \rangle$. The group $\Omega^{\pm 1}(X)$ acts on $\mathrm{E}\Pi\mathrm{A}_n$ in the natural way, by permuting and/or inverting the elementary palindromic automorphisms. A defining set of relations for $\mathrm{E}\Pi\mathrm{A}_n$ is given by

1. $[P_{ij}, P_{ik}] = 1$,
2. $[P_{ij}, P_{kl}] = 1$, and
3. $P_{ij}P_{jk}P_{ik} = {P_{ik}}^{-1}P_{jk}P_{ij}$,

where i, j, k, l are pairwise distinct. Note that for $n = 2$ or 3, some of these relations are not defined. Removing undefined relations from the list gives a complete set of defining relations in these cases.

A striking comparison is made by Collins between these defining relators for $\mathrm{E}\Pi\mathrm{A}_n$, and a finite presentation for the *pure symmetric automorphism group of* F_n, denoted $\mathrm{P}\Sigma A_n$, of automorphisms that take each $x \in X$ to a conjugate of itself. Let $C_{ij} \in \mathrm{P}\Sigma A_n$ map x_i to ${x_j}^{-1}x_ix_j$ $(i \neq j)$ and fix all $x_k \in X$ with $k \neq i$. Then $\mathrm{P}\Sigma A_n$ is generated by the set $\{C_{ij} \mid i \neq j\}$, subject to the defining relations

1. $[C_{ij}, C_{ik}] = 1$,
2. $[C_{ij}, C_{kl}] = 1$, and
3. $C_{ij}C_{jk}C_{ik} = C_{ik}C_{jk}C_{ij}$,

where i, j, k, l are pairwise distinct. Note that these abstractly differ from the relations defining $\mathrm{E}\Pi\mathrm{A}_n$ only in the exponent of C_{ik} in relations of type 3. This comparison of Collins motivated the suggestion that $\mathrm{E}\Pi\mathrm{A}_n$ could be understood by adapting methods that had been used to analyse $\mathrm{P}\Sigma A_n$. Indeed, this proved fruitful, with Glover-Jensen [31] letting $\mathrm{E}\Pi\mathrm{A}_n$ act on a contractible subcomplex of *auter space* (an analogue of Teichmüller space for $\mathrm{Aut}(F_n)$) to study torsion and cohomological properties of $\Pi\mathrm{A}_n$.

Using graph folding techniques of Stallings, we obtain a new proof of finite generation of $\Pi\mathrm{A}_n$, as well as finding generating sets for certain fixed point subgroups of $\Pi\mathrm{A}_n$.

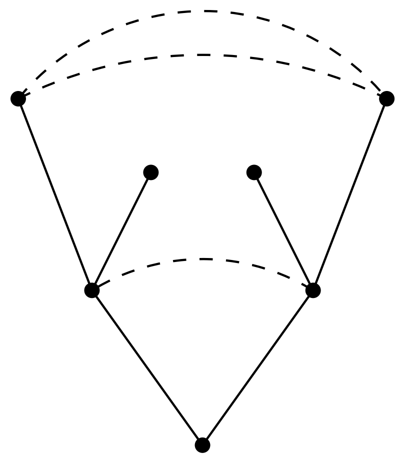

Figure 2.9: *An example of an arch, with base point v. The dashed edges indicate the bridges that have been added to the trees that were glued together at the base point.*

Proposition 2.2.8. *Fix $0 \leq k \leq n$, and let $\Pi\mathrm{A}_n(k)$ consist of automorphisms which fix $x_1, \dots, x_k$, with the convention that $\Pi\mathrm{A}_n(0) = \Pi\mathrm{A}_n$. A finite generating set for $\Pi\mathrm{A}_n(k)$ is*

$$\left[\Omega^{\pm 1}(X) \cap \Pi\mathrm{A}_n(k)\right] \cup \{P_{ij} \mid i > k\}.$$

Proof. The idea behind this proof was inspired by a proof of Wade [58, Theorem 4.1]. We begin by introducing some terminology. Let $\phi : S \to T$ be an isomorphism of finite trees. For a vertex (resp. edge) r of S, denote by r' the image of r under ϕ. Choose a distinguished vertex v of S, of valence 1. An *arch of S at v* (see Figure 2.9) is the graph formed by gluing S to T along v and v', then for each vertex $r \in S \setminus \{v\}$, adding some (possibly zero) number of edges between r and r'. We refer to these new edges as *bridges*. The image of v in the arch forms a natural base point, and any edge with v as one of its endpoints is called a *stem*. By an *wedge of arches* we mean a collection of arches glued together at their base points.

Let $\alpha \in \Pi\mathrm{A}_n(k)$ and let R_n be n copies of S^1 glued together at a single point, where each S^1 is endowed with an orientation to give a canonical generating set for $\pi_1(R_n) = F_n$. We may realise α as a map of graphs $\theta : Y \to R_n$, where Y is the result of subdividing each S^1 of R_n into the appropriate number of edges, and 'spelling out' the word $\alpha(x_i)$ on the ith copy of S^1: precisely, the jth edge of the oriented, subdivided S^1 corresponding to $\alpha(x_i)$ is mapped to the loop in R_n corresponding to the jth letter of $\alpha(x_i)$, correctly oriented. We

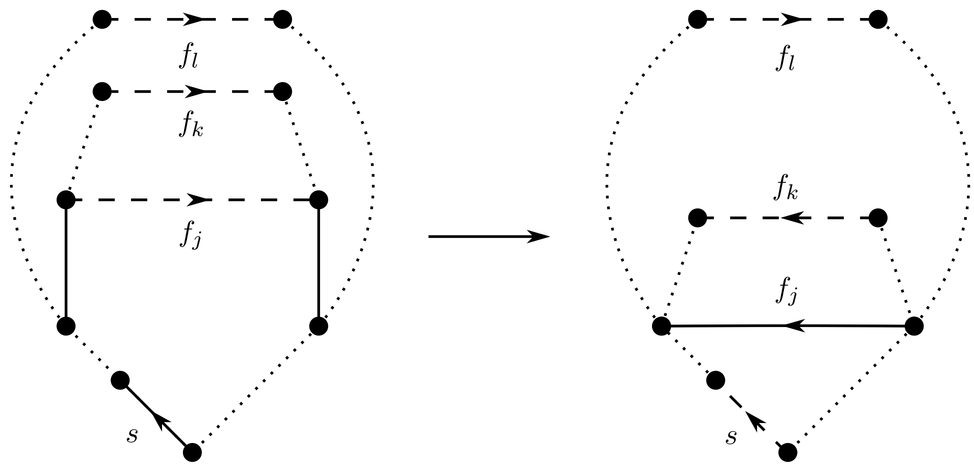

Figure 2.10: *The two adjacent solid edges are folded onto f_j. The dashed edges represent edges excluded from the graph's chosen maximal tree. In order to record what effect this type B 2-fold has on the branded graph's associated automorphism, we must swap f_j into the maximal tree, in place of the stem s.*

now use graph folding to write α as a product of permutations, inversions and elementary palindromic automorphisms.

We use the terminology of Wade [58], which we introduced in Section 2.2.3. Observe that Y is a wedge of n arches, each of which arises from an isomorphism of trees $\phi_i : S_i \to T_i$ $(1 \leq i \leq n)$. Due to the symmetry of a palindromic word, folds come together in natural pairs. Consider folds of type 1. For instance, if we are able to fold together two edges $h_i \in S_i$ and $h_j \in S_j$, since $\theta(h_i) = \theta(h_j)$, then we will also be able to fold together $\phi_i(h_i)$ and $\phi_j(h_j)$, as they will also both have the same image under θ. We call this pair of folds a *type A 2-fold*. We may also have a sequence of edges (h_{j-1}, h_j, h_{j+1}) in S_i mapped under θ to the sequence $(\bar{x}, x, \bar{x})$ where h_j is a bridge and x is some edge in R_n. We fold h_{j-1} and h_{j+1} onto h_j, and call this pair of folds a *type B 2-fold*. Such a fold is seen in Figure 2.10. Doing either of these 2-folds to Y yields another, different wedge of arches. The argument just used also applies to this new wedge of arches, and so we may continue to carry out 2-folds, each of which reduces the number of edges in the graph.

In order to see what effect these 2-folds have on $\alpha \in \Pi\mathrm{A}_n(k)$, we must keep track of a canonical maximal tree T we define on Y. The edges of Y not in T are the bridges coming from each arch. In order to carry out a type B 2-fold we must swap the bridge f_j into the maximal tree. Recall $p_{i(f_j)}$ is the unique reduced path in T joining the base point to the initial vertex of f_j. Apart from one degenerate case, which we deal with separately, we may always swap f_j into the maximal tree T by excluding the stem appearing in $p_{i(f_j)}$. We show that the result of swapping maximal trees, doing a type B 2-fold, then swapping back to the maximal tree where all bridges are excluded is to carry out an elementary palindromic automorphism, $P_{ij}^{\epsilon_k}$, to some members of X.

Let $\hat{\theta} : \hat{Y} \to R_n$ be a map of graphs obtained by carrying out a sequence of 2-folds to the map $\theta : Y \to R_n$. With our canonical maximal tree T, these data constitute a branding $\mathcal{H}_1$ of $\hat{Y}$. Suppose that we wish to do a type B 2-fold onto the bridge f_j, as seen in Figure 2.10. First we swap f_j into the maximal tree in place of the stem s in $p_{i(f_j)}$ to produce a new branding $\mathcal{H}_2$ of $\hat{Y}$. Then by Proposition 2.2.7, we have

$$B_{\mathcal{H}_1} = B_{\mathcal{H}_2} \cdot (A, {x_j}^{\epsilon}),$$

where A consists precisely of the elements ${x_k}^{\epsilon_k}$ when $p_{i(f_k)}$ or $p_{t(f_k)}$ involve the edge s (with ϵ_k chosen to be 1 or -1 accordingly). We then fold the two edges onto the bridge, and obtain a new graph Y' with branding $\mathcal{H}_3$. By Proposition 2.2.5, we have $B_{\mathcal{H}_2} = B_{\mathcal{H}_3}$. Finally, we return to the canonical maximal tree of Y' by swapping s back into the tree. As per the instructions in Section 2.2.3, we do this by excluding the edge $\overline{f_j}$, and obtain a branding $\mathcal{H}_4$. Again by Proposition 2.2.7, we see that

$$B_{\mathcal{H}_3} = B_{\mathcal{H}_4} \cdot W,$$

for some Whitehead automorphism W.

It is straightforward to verify what the automorphism $W \cdot (A, x_j)$ does to the members of the free basis X. Let $x_l \in X$ be such that f_l is as shown in Figure 2.10 (that is, $i(f_j) \notin p_{i(f_l)}$). Then $W \cdot (A, x_j)$ fixes x_l. Let $x_k \in X$ be such that f_k is as shown in Figure 2.10 (that is, $i(f_k) \in p_{i(f_k)}$). Then $W \cdot (A, x_j)$ maps x_k to ${x_j}^{\epsilon_k} x_k {x_j}^{\epsilon_k}$, where ϵ_k depends on the orientations in the graph $\hat{Y}$.

The only degenerate case of the above is when one (and hence both) of the edges we want to fold onto a bridge is a stem. In this case, we change maximal trees as before then fold one of the stems onto the bridge with a type 1 fold. This causes the other stem to become a loop, around which we fold the bridge using a type 2 fold. The Whitehead automorphisms associated to these three steps compose as before to give a product of elementary palindromic automorphisms.

Carrying out a sequence of 2-folds of types A and B eventually produces a map $R_n \to R_n$, and so we complete the algorithm by applying the appropriate Whitehead automorphism of type 1. Notice that since $\alpha \in \Pi A_n(k)$, the graph Y we constructed has a single loop at the base point for each x_i $(1 \leq i \leq k)$, as $\alpha(x_i) = x_i$, so the first k ordered loops of R_n were not subdivided to form Y. Thus, while folding such a graph Y, we only need Collins' generators (the elementary palindromic automorphisms and members of $\Omega^{\pm 1}(X)$) that fix the first k members of the free basis X. The proposition is thus proved. □

Corollary 2.2.9. *The group* $\mathrm{P}\Pi\mathrm{A}_n(k)$ *of pure palindromic automorphisms fixing* $x_1, \ldots, x_k$ $(0 \leq k \leq n)$ *is generated by the set* $\{P_{ij},\, \iota_i \mid i > k\}$.

2.2.5 The level 2 congruence subgroup of $\mathrm{GL}(n, \mathbb{Z})$

Let $\Gamma_n[2]$ denote the kernel of the map $\mathrm{GL}(n, \mathbb{Z}) \to \mathrm{GL}(n, \mathbb{Z}_2)$ given by reducing matrix entries mod 2. This is the so-called *principal level 2 congruence subgroup* of $\mathrm{GL}(n, \mathbb{Z})$. Let S_{ij} be the matrix with 1s on the diagonal, 2 in the (i, j) position and 0s elsewhere, and let O_i be the matrix which differs from the identity matrix only in having -1 in the (i, i) position. The following lemma verifies a well-known generating set for $\Gamma_n[2]$ (see, for example, McCarthy-Pinkall [46, Corollary 2.3].

Lemma 2.2.10. *The set* $\{O_i, S_{ij} \mid 1 \leq i \neq j \leq n\}$ *generates* $\Gamma_n[2]$.

Proof. Observe that we may think of the matrices S_{ij} as corresponding to carrying out 'even' row operations: that is, adding an even multiple of one matrix row to another. Let u be the first column of some matrix in $\Gamma_n[2]$, and denote by $u^{(i)}$ the ith entry of u. Let v_1 be the standard column vector with 1 in the first entry and 0s elsewhere.

Claim: The column u can be reduced to $\pm v_1$ using even row operations.

We use induction on $|u^{(1)}|$. For $|u^{(1)}| = 1$, the result is trivial. Now suppose $|u^{(1)}| > 1$. As in the proof of Proposition 2.2.2, we deduce that there must be some $u^{(j)}$ which is not a multiple of $u^{(1)}$. By the Division Algorithm, there exist $q, r \in \mathbb{Z}$ such that $u^{(j)} = q|u^{(1)}| + r$, with $0 \leq r < |u^{(1)}|$. If q is not even, we instead write $u^{(j)} = (q+1)|u^{(1)}| + (r - |u^{(1)}|)$. Note that if q is odd, then $r \neq 0$, since $u^{(1)}$ is odd and $u^{(j)}$ is even, and so $-|u^{(1)}| < r - |u^{(1)}|$. Depending on the parity of q, we do the appropriate number of even row operations to replace $u^{(j)}$ with r or $r - |u^{(1)}|$. In both cases, we have replaced $u^{(j)}$ with an integer of absolute value smaller than $|u^{(1)}|$. It is clear that now we may reduce the absolute value of $u^{(1)}$ by either adding or subtracting twice the jth row from the first row, and so by induction we have proved the claim.

We now induct on n to prove the lemma. It is clear that $\Gamma_1[2] = \langle O_1 \rangle$. Using the above claim, we may assume that we have reduced $M \in \Gamma_n[2]$ so it is of the form

$$\left[\begin{array}{c|c} \pm 1 & * \\ \hline 0 & N \end{array}\right],$$

where $N \in \Gamma_{n-1}[2]$. Our aim is to further reduce M to the identity matrix using the set of matrices in the statement of the lemma. By induction, we may assume that N can be reduced to the identity matrix using the appropriate members of $\{E_{ij}, O_i \mid i, j > 1\}$. Then we simply use even row operations to fix the top row, and finish by applying O_1 if necessary. $\square$

By Lemma 2.2.10, the restriction of the short exact sequence

$$1 \longrightarrow \mathrm{IA}_n \longrightarrow \mathrm{Aut}(F_n) \longrightarrow \mathrm{GL}(n, \mathbb{Z}) \longrightarrow 1$$

to $\mathrm{P}\Pi\mathrm{A}_n$ gives the short exact sequence

$$1 \longrightarrow \mathcal{PI}_n \longrightarrow \mathrm{P}\Pi\mathrm{A}_n \longrightarrow \Gamma_n[2] \longrightarrow 1,$$

since P_{ij} maps to S_{ji} and ι_i maps to O_i.

The rest of this chapter is concerned with finding a generating set for $\mathcal{PI}_n$. We find such a set by constructing a new complex on which $\mathcal{PI}_n$ acts in a suitable way. We then apply a theorem of Armstrong [2] to conclude that $\mathcal{PI}_n$ is generated by the action's vertex stabilisers. In the following section, we define the complex and use it to prove Theorem 2.1.1.

2.3 The complex of partial π-bases

Day-Putman [24] use the *complex of partial bases* of F_n, denoted $\mathcal{B}_n$, to derive a generating set for IA_n. We build a complex modelled on $\mathcal{B}_n$, and follow the approach of Day-Putman to find a generating set for $\mathcal{PI}_n$.

Fix $X := \{x_1, \ldots, x_n\}$ as a free basis of F_n. A *π-basis* is a set of palindromes on X which also forms a free basis of F_n. A *partial π-basis* is a set of palindromes on X which may be extended to a π-basis. The *complex of partial π-bases* of F_n, denoted $\mathfrak{B}_n^\pi$, is defined to be the simplicial complex whose $(k-1)$-simplices correspond to partial π-bases $\{w_1, \ldots, w_k\}$. We postpone until Section 2.4 the proof of the following theorem on the connectedness of $\mathfrak{B}_n^\pi$.

Theorem 2.3.1. *For $n \geq 3$, the complex $\mathfrak{B}_n^\pi$ is simply-connected.*

Our complex $\mathfrak{B}_n^\pi$ is technically not a subcomplex of $\mathcal{B}_n$, as the vertices of $\mathcal{B}_n$ are taken to be conjugacy classes, rather than genuine members of F_n. We ignore this technicality, as Proposition 2.2.1 shows that if two odd-length palindromes are conjugate, they are equal. It is clear, however, that $\mathfrak{B}_n^\pi$ is isomorphic to a subcomplex of $\mathcal{B}_n$.

There is an obvious simplicial action of $\Pi\mathrm{A}_n$ on $\mathfrak{B}_n^\pi$. This action is, by definition, transitive on the set of k-simplices, for each $0 \leq k < n$. Further, $\mathcal{PI}_n$ acts without rotations: that is, if $\phi \in \mathcal{PI}_n$ stabilises a simplex σ of $\mathfrak{B}_n^\pi$, then it fixes σ pointwise. The quotient of $\mathfrak{B}_n^\pi$ by $\mathcal{PI}_n$ is highly connected, by a theorem of Charney [14].

Theorem 2.3.2 (Charney)**.** *For $n \geq 3$, the quotient $\mathfrak{B}_n^\pi/\mathcal{PI}_n$ is $(n-3)$-connected.*

The proof of this theorem is discussed in Section 2.4.

These theorems allow us to apply the following theorem of Armstrong [2] to the action of $\mathcal{PI}_n$ on $\mathfrak{B}_n^\pi$, for $n \geq 4$. The statement of the theorem is as given in Day-Putman [24].

Theorem 2.3.3. *Let G act simplicially on a simply-connected simplicial complex X, without rotations. Then G is generated by the vertex stabilisers of the action if and only if X/G is simply-connected.*

We analyse the vertex stabilisers of $\mathcal{PI}_n$ using an inductive argument. It is known that $\mathcal{PI}_1 = 1$ and $\mathcal{PI}_2 = 1$: the latter inequality follows from the fact that $\mathrm{IA}_2 = \mathrm{Inn}(F_2)$ [49] and $\mathrm{Inn}(F_n) \cap \Pi\mathrm{A}_n = 1$ for $n \geq 1$. We treat the $n = 3$ case differently, as the quotient $\mathfrak{B}_3^\pi/\mathcal{PI}_3$ is not simply-connected, and so does not allow us to apply Armstrong's theorem directly. This treatment is postponed until Section 2.5.

2.3.1 A Birman exact sequence

We require a version of the free group analogue of the Birman exact sequence, as developed by Day-Putman [23]. Recall that $\mathrm{P\Pi A}_n(k)$ consists of the pure palindromic automorphisms fixing $x_1, \ldots, x_k$.

Proposition 2.3.4. *For $0 \leq k \leq n$, there exists the split short exact sequence*

$$1 \longrightarrow \mathcal{J}_n(k) \longrightarrow \mathrm{P\Pi A}_n(k) \longrightarrow \mathrm{P\Pi A}_{n-k} \longrightarrow 1,$$

where $\mathcal{J}_n(k)$ is the normal closure in $\mathrm{P\Pi A}_n(k)$ of the set $\{P_{ij} \mid i > k,\ j \leq k\}$.

Proof. A map $\mathrm{P\Pi A}_n(k) \to \mathrm{P\Pi A}_{n-k}$ is induced by the map $F_n \to F_{n-k}$ that trivialises $x_1, \ldots, x_k$. The existence of the split short exact sequence follows from Corollary 2.2.9. $\square$

Our 'Birman kernel' $\mathcal{J}_n(k)$ is rather worse behaved than the analogous Birman kernel of Day-Putman: theirs, denoted $\mathcal{K}_{n,k,l}$, is finitely generated whereas it may be shown by adapting the proof of their Theorem E [23] that $\mathcal{J}_n(k)$ is not. This difference is due in part to the fact that their version of $\mathrm{P\Pi A}_n(k)$ need only fix each of $x_1, \ldots, x_k$ up to conjugacy. The lack of finite generation of $\mathcal{J}_n(k)$ is, however, not an obstacle to the goal of this chapter: we only require that $\mathcal{J}_n(k)$ is *normally* generated by a finite set.

Our Birman exact sequence projects into $\mathrm{GL}(n, \mathbb{Z})$ in an obvious way, made precise in the following lemma. Let v_i denote the image of $x_i \in F_n$ under the abelianisation map. We denote by $\Gamma_n[2](k)$ the members of $\Gamma_n[2]$ which fix $v_1, \ldots, v_k \in \mathbb{Z}^n$, and by $\mathcal{H}_n(k)$ the group $\mathrm{Hom}(\mathbb{Z}^{n-k}, (2\mathbb{Z})^k)$.

Lemma 2.3.5. *Fix $0 \leq k \leq n$. Then there exists the following commutative diagram of split short exact sequences*

$$\begin{tikzcd}
1 \arrow[r] & \mathcal{J}_n(k) \arrow[r] \arrow[d, two heads] & \mathrm{P\Pi A}_n(k) \arrow[r] \arrow[d, two heads] & \mathrm{P\Pi A}_{n-k} \arrow[r] \arrow[d, two heads] \arrow[l, bend left, "s"] & 1, \\
1 \arrow[r] & \mathcal{H}_n(k) \arrow[r] & \Gamma_n[2](k) \arrow[r] & \Gamma_{n-k}[2] \arrow[r] \arrow[l, bend left, "t"] & 1
\end{tikzcd}$$

where s and t are the obvious splitting homomorphisms.

Proof. The top row is given by Proposition 2.3.4. A generating set for $\Gamma_n[2](k)$ follows from the proof of Lemma 2.2.10; it is precisely the image in $\mathrm{GL}(n,\mathbb{Z})$ of $\{P_{ij},\ \iota_i \mid i > k\}$, the generating set of $\mathrm{P\Pi A}_n(k)$ given by Corollary 2.2.9. The bottom row then follows by an argument similar to the proof of Proposition 2.3.4, noting that the kernel is generated by the images of P_{ij} $(i > k,\ j \leq k)$. It is straightforward to verify that this kernel is $\mathrm{Hom}(\mathbb{Z}^{n-k},(2\mathbb{Z})^k)$. Intuitively, $\alpha \in \mathrm{Hom}(\mathbb{Z}^{n-k},(2\mathbb{Z})^k)$ is encoding how many (even) multiples of v_j $(1 \leq i \leq k)$ are added to each v_i $(k < j \leq n)$.

The right hand vertical map follows from Lemma 2.2.10. It is clear that all the arrows commute, and that the splitting homomorphisms s and t are compatible with the commutative diagram, so the proof is complete. □

2.3.2 A generating set for $\mathcal{J}_n(1) \cap \mathcal{PI}_n$

By mapping $\mathrm{P\Pi A}_n(k)$ into $\Gamma_n[2](k)$ then conjugating the normal subgroup $\mathcal{H}_n(k)$, we obtain a homomorphism $\alpha_k : \mathrm{P\Pi A}_n(k) \to \mathrm{Aut}(\mathcal{H}_n(k))$. Setting $k = 1$, we obtain the following lemma.

Lemma 2.3.6. *The group $\mathcal{J}_n(1) \cap \mathcal{PI}_n$ is normally generated in $\mathcal{J}_n(1)$ by the set*

$$\{[P_{ij}, P_{i1}],\ [P_{ij}, P_{j1}]P_{i1}^2 \mid 1 < i \neq j \leq n\}.$$

Proof. By Lemma 2.3.5, there is a short exact sequence

$$1 \longrightarrow \mathcal{J}_n(1) \cap \mathcal{PI}_n \longrightarrow \mathcal{J}_n(1) \longrightarrow \mathcal{H}_n(1) \longrightarrow 1.$$

The set $Y := \{\phi P_{j1} \phi^{-1} \mid \phi \in \mathrm{P\Pi A}_n(1), 1 < j \leq n\}$ generates $\mathcal{J}_n(1)$ by Proposition 2.3.4. Let a_j denote the image of P_{j1} in $\mathrm{GL}(n,\mathbb{Z})$. A straightforward calculation verifies that the set $\{a_j\}$ is a free abelian basis for $\mathcal{H}_n(1)$; this follows since $\mathcal{H}_n(1) = \langle S_{1k} \rangle$ $(k > 1)$, with this generating set being a free abelian basis for $\mathcal{H}_n(1)$.

For $\phi \in \mathrm{P\Pi A}_n(1)$, let $\bar{\phi}$ denote the image of ϕ in $\Gamma_n[2](1)$, and let $\bar{Y}$ denote the image of Y. The set of relations

$$\{[a_i, a_j] = 1,\ \bar{\phi} a_i \bar{\phi}^{-1} = \alpha_1(\phi)(a_i) \mid 1 < i \neq j \leq n,\ \phi \in \mathrm{P\Pi A}_n(1)\},$$

together with the generating set $\bar{Y}$, form a presentation for $\mathcal{H}_n(k)$. It is clear that the image of any member of Y in $\mathcal{H}_n(1)$ is a word on the free abelian basis $\{a_i\}$, and that this word is determined by the homomorphism α_1.

It is a standard fact (see, for example, the proof of Theorem 2.1 in Magnus-Karrass-Solitar [44]) that $\mathcal{J}_n(1) \cap \mathcal{PI}_n$ is normally generated in $\mathcal{J}_n(1)$ by the obvious lifts of the (infinitely many) relators in the given presentation for $\mathcal{H}_n(1)$. The relators of the form $[a_i, a_j]$ have trivial lift, and so are not required in the generating set. Let C be the finite generating set for $\mathrm{P\Pi A}_n(1)$ given by Corollary 2.2.9. It can be shown that the obvious lift of the finite set of relators

$$D := \{\bar{c} a_j \bar{c}^{-1} \alpha_1(c)(a_j)^{-1} \mid c \in C, 1 < j \leq n\}$$

suffices to normally generate $\mathcal{J}_n(1) \cap \mathcal{PI}_n$ in $\mathcal{J}_n(1)$. This may be seen using a simple induction argument on the length of the word $\phi \in \mathrm{P\Pi A}_n(1)$ on C, as we now show. Let $\phi = c_1 \dots c_k$ with $c_i \in C^{\pm 1}$. We wish to show that

$$\mathcal{W} := \left[\bar{c}_1 \dots \bar{c}_k a_j \bar{c}_k^{-1} \dots \bar{c}_1^{-1}\right] \alpha_1(c_1 \dots c_k)(a_j)^{-1}$$

lies in the normal closure of the lift of D in $\mathcal{J}_n(1)$. By definition,

$$\alpha_1(c_1 \dots c_k)(a_j)^{-1} = \alpha_1(c_1)\left[\alpha_1(c_2 \dots c_k)(a_j)^{-1}\right] = \bar{c}_1 \cdot \alpha_1(c_2 \dots c_k)(a_j)^{-1} \cdot \bar{c}_1^{-1},$$

and so

$$\mathcal{W} = \left[\bar{c}_1(c_2 \dots c_k a_j \bar{c}_k^{-1} \dots \bar{c}_2^{-1})\bar{c}_1^{-1}\right]\left[\bar{c}_1 \alpha_1(c_2 \dots c_k)(a_j)^{-1} \bar{c}_1^{-1}\right].$$

Induction now allows us to conclude that $\mathcal{W}$ lies in the normal closure, as desired.

All that remains is to verify that the obvious lift of D is the set given in the statement of the lemma; this is a straightforward calculation, which is summarised in Table 2.1.

□

2.3.3 Proof of Theorem 2.1.1

We now prove Theorem 2.1.1, using the action of $\mathcal{PI}_n$ on $\mathfrak{B}_n^\pi$.

Proof of Theorem 2.1.1. The action of $\mathcal{PI}_n$ on $\mathfrak{B}_n^\pi$ is simplicial and without rotations. Combining Theorems 2.3.1, 2.3.2 and 2.3.3, we conclude that for $n \geq 4$, $\mathcal{PI}_n$ is generated by the vertex stabilisers of the action on $\mathfrak{B}_n^\pi$.

Generator $c \in C$	The lift of $\bar{c}a_j\bar{c}^{-1}\alpha_1(c)(a_j)^{-1}$ to $\mathcal{J}_n(1)$
ι_j	1
ι_i	1
P_{k1}	1
P_{il}	1
P_{ij}	$[P_{ij}, P_{j1}]{P_{i1}}^2$
P_{ji}	$[P_{ji}, P_{j1}]$

Table 2.1: *The lifts of the members of the set D, where the indices i, j, k and l are taken to be pairwise distinct, with $i, j, k \neq 1$.*

Recall that $\mathcal{PI}_n(1)$ denotes the stabiliser of the vertex x_1. Since $\Pi\mathrm{A}_n$ acts transitively on the vertices of $\mathfrak{B}_n^\pi$, the stabiliser in $\mathcal{PI}_n$ of any vertex is conjugate in $\Pi\mathrm{A}_n$ to $\mathcal{PI}_n(1)$. Lemma 2.3.5 gives us the split short exact sequence

$$1 \longrightarrow \mathcal{J}_n(1) \cap \mathcal{PI}_n \longrightarrow \mathcal{PI}_n(1) \longrightarrow \mathcal{PI}_{n-1} \longrightarrow 1.$$

We induct on n. By the above split short exact sequence, to generate $\mathcal{PI}_n(1)$ it suffices to combine a generating set of $\mathcal{J}_n(1) \cap \mathcal{PI}_n(1)$ with a lift of one of $\mathcal{PI}_{n-1}$.

We begin with the base case, $n = 3$. In Section 2.5, we verify that the presentation of $\Gamma_3[2]$ given in Corollary 2.1.2 is correct when $n = 3$. Given the short exact sequence

$$1 \longrightarrow \mathcal{PI}_3 \longrightarrow \mathrm{P}\Pi\mathrm{A}_3 \longrightarrow \Gamma_3[2] \longrightarrow 1,$$

we may take the obvious lifts of the relators in this presentation as a normal generating set for $\mathcal{PI}_3$ in $\mathrm{P}\Pi\mathrm{A}_3$. Relators 1-6 are trivial when lifted, while relator 8 and 9 lift to $[P_{ij}, P_{ik}]$ and (an automorphism equal to) $P_{ik}[P_{ij}, P_{ik}]{P_{ik}}^{-1}$, respectively, both of which are conjugate to $[P_{12}, P_{13}]$. Finally, relator 10 lifts to

$$(P_{23}{P_{13}}^{-1}P_{31}P_{32}P_{12}{P_{21}}^{-1})^2,$$

so the base case $n = 3$ is true.

Now suppose $n > 3$. By induction, the group $\mathcal{PI}_{n-1}$ is normally generated by $[P_{42}, P_{43}]$ and $(P_{23}{P_{43}}^{-1}P_{34}P_{32}P_{42}{P_{24}}^{-1})^2$, say, in $\Pi\mathrm{A}_{n-1}$. We lift this normal generating set to $\mathcal{PI}_n(1)$ in the obvious way.

By Lemma 2.3.6, we need only add in $\mathcal{J}_n(1)$-conjugates of the words $[P_{ij}, P_{i1}]$ and $[P_{ij}, P_{j1}]P_{i1}^2$, for $1 < i \neq j \leq n$. The former are clearly conjugate in $\Pi\mathrm{A}_n$ to $[P_{12}, P_{13}]$. For the latter, observe that

$$[P_{ij}, P_{j1}]P_{i1}^2 = [P_{ij}, P_{i1}^{-1}],$$

which again is conjugate to $[P_{12}, P_{13}]$, so we are done. □

Corollary 2.1.2 follows immediately from Theorem 2.1.1. Since $\Gamma_n[2] \cong \mathrm{P}\Pi\mathrm{A}_n/\mathcal{PI}_n$, by adding the normal generators in Theorem 2.1.1 (and all words obtained by permuting their indices) as relators to Collins' presentation of $\mathrm{P}\Pi\mathrm{A}_n$, we obtain a finite presentation of $\Gamma_n[2]$. Applying the obvious Tietze transformations yields the presentation given in Corollary 2.1.2.

We end this section by proving that the (normal) generator $(P_{23}{P_{13}}^{-1}P_{31}P_{32}P_{12}{P_{21}}^{-1})^2$ in the statement of Theorem 2.1.1 is necessary.

Proposition 2.3.7. *For $n \geq 3$, the group normally generated by $[P_{12}, P_{13}]$ in $\Pi\mathrm{A}_n$ is a proper subgroup of $\mathcal{PI}_n$.*

Proof. Suppose $\mathcal{PI}_n$ is the normal closure of $[P_{12}, P_{13}]$ in $\Pi\mathrm{A}_n$. Then the orbit of $[P_{12}, P_{13}]$ under the action of the symmetric group on the free basis X produces a normal generating set for $\mathcal{PI}_n$ in $\mathrm{P}\Pi\mathrm{A}_n$. Adding these to the presentation of $\mathrm{P}\Pi\mathrm{A}_n$ as relators yields a finite presentation $\mathcal{Q}$ of $\Gamma_n[2]$, which may be altered using Tietze transformations so that it looks like the presentation in Corollary 2.1.2, with relator 10 removed.

We know that

$$\chi := (S_{32}{S_{31}}^{-1}S_{13}S_{23}S_{21}{S_{12}}^{-1})^2$$

is trivial in $\Gamma_n[2]$, and so we should be able to deduce this as a consequence of the relations in $\mathcal{Q}$. We derive a contradiction by showing that χ is non-trivial in the group presented by $\mathcal{Q}$. Observe that by killing all the generators of $\Gamma_n[2]$ except for S_{12} and S_{21}, we surject onto the free Coxeter group generated by the images of S_{12} and S_{21}, say A and B, respectively. This is easily verified by examining the relators of $\mathcal{Q}$. The image of χ under this map is $ABAB \neq 1$, and so χ is non-trivial in the group presented by $\mathcal{Q}$. Therefore the normal closure of $[P_{12}, P_{13}]$ in $\Pi\mathrm{A}_n$ is not all of $\mathcal{PI}_n$. □

Note that in the proof of Proposition 2.3.7 we also showed that relators 1–9 of Corollary 2.1.2 are not a sufficient set of relators that hold between the O_i and S_{jk}: relator 10 is not a consequence of the others.

Corollary 2.3.8. *The complex $\mathfrak{B}_3^\pi/\mathcal{PI}_3$ is not simply-connected.*

Proof. By Theorem 2.3.3, the complex $\mathfrak{B}_3^\pi/\mathcal{PI}_3$ is simply-connected if and only if $\mathcal{PI}_3$ is generated by the vertex stabilisers of the action. Proposition 2.3.7 shows that vertex stabilisers do not suffice to generate $\mathcal{PI}_3$, so the quotient is not simply-connected. □

2.4 The connectivity of $\mathfrak{B}_n^\pi$ and its quotient

In this section, we determine the levels of connectivity of $\mathfrak{B}_n^\pi$ and $\mathfrak{B}_n^\pi/\mathcal{PI}_n$. The former is found to be simply-connected, following the same approach as Day-Putman [24], while the latter is shown to be closely related to a complex already studied by Charney [14], which is $(n-3)$-connected.

2.4.1 The connectivity of $\mathfrak{B}_n^\pi$

First, we recall the definition of the Cayley graph of a group. Let G be a group with finite generating set S. The *Cayley graph of G with respect to S*, denoted $\mathrm{Cay}(G,S)$, is the graph with vertex set G and edge set $\{(g,gs) \mid g \in G, s \in S^{\pm 1}\}$, where an ordered pair (x,y) indicates that vertices x and y are joined by an edge. If $s \in S$ has order 2, we identify each pair of edges (g,gs) and (g,gs^{-1}) for each $g \in G$, to ensure that the Cayley graph is simplicial.

We establish Theorem 2.3.1 by constructing a map Ψ from the Cayley graph of $\Pi\mathrm{A}_n$ to $\mathfrak{B}_n^\pi$ and demonstrating that the induced map of fundamental groups is both surjective and trivial. We require the Cayley graph of $\Pi\mathrm{A}_n$ with respect to a particular generating set, which we now describe. Assume that $n \geq 3$. For $1 \leq i \neq j < n$, let t_{ij} permute x_i and x_j, fixing x_k with $k \neq i,j$. Using the symmetric group action on X, we deduce from Proposition 2.2.8 that we may generate $\Pi\mathrm{A}_n$ using the set

$$Z := \{t_{ij}, \iota_2, \iota_3, P_{21}, P_{23}, P_{31}, P_{34} \mid 1 \leq i \neq j \leq n\}.$$

We may use the symmetric group action on X to streamline the presentation of ΠA_n given in Section 2.2, to obtain the following list of defining relators for ΠA_n on the generating set Z:

1. $t_{ij} = t_{ji}$,
2. $t_{ij}^2 = 1$,
3. $ut_{ij}u^{-1} = t_{u(i)u(j)}$,
4. ${\iota_2}^2 = 1$,
5. $(\iota_2\iota_3)^2 = 1$,
6. $[\iota_2, P_{31}] = 1$,
7. $(\iota_2 P_{21})^2 = 1$,
8. $(\iota_3 P_{23})^2 = 1$,
9. $P_{23}P_{31}P_{21} = P_{21}^{-1}P_{31}P_{23}$,
10. $[P_{21}, P_{31}] = 1$,
11. $[P_{21}, P_{34}] = 1$,
12. $\iota_3 = t_{23}\iota_2 t_{23}$,
13. $P_{31} = t_{23}P_{21}t_{23}$,
14. $P_{23} = t_{13}P_{21}t_{13}$,
15. $P_{34} = t_{14}t_{23}P_{21}t_{23}t_{14}$,
16. $P_{21} = wP_{21}w^{-1}$ for $w \in \mathcal{W}$,
17. $\iota_2 = v\iota_2 v^{-1}$ for $v \in \mathcal{V}$,

where $1 \leq i \neq j \leq n$, $u \in \{t_{ij}\}$ and $\mathcal{W}$ and $\mathcal{V}$ are the sets of words on $\{t_{ij}\}$ that fix x_1 and x_2, and only x_2, respectively. The relations of type 16 and 17 arise due to the streamlining of the presentation of $\Pi A_n = \mathrm{E}\Pi A_n \rtimes \Omega^{\pm 1}(X)$ given in Section 2.2. Note that relations $1 - 3$ are a complete set of relations for the symmetric group, when generated by the transpositions $\{t_{ij}\}$ [52].

We now consider the Cayley graph $\mathrm{Cay}(\Pi A_n, Z)$. Observe that for each $z \in Z^{\pm 1}$, either $z(x_1) = x_1$ or $\{x_1, z(x_1)\}$ forms a partial π-basis for F_n. This allows us to construct a map of complexes from the star of the vertex 1 in $\mathrm{Cay}(\Pi A_n, Z)$ to $\mathfrak{B}_n^\pi$, by mapping an edge $z \in Z^{\pm 1}$ to the edge $v_1 - z(v_1)$ (which may be degenerate). Using the actions of ΠA_n on $\mathrm{Cay}(\Pi A_n, Z)$ and $\mathfrak{B}_n^\pi$, we can extend this map to a map of complexes $\Psi : \mathrm{Cay}(\Pi A_n, Z) \to \mathfrak{B}_n^\pi$. Explicitly, Ψ takes a vertex $z_1 \dots z_r$ of $\mathrm{Cay}(\Pi A_n, Z)$ ($z_i \in Z^{\pm 1}$) to the vertex $z_1 \dots z_r(x_1)$.

Proof of Theorem 2.3.1. This proof is modelled on Day-Putman's proof of their Theorem A [24]. Let

$$\Psi_* : \pi_1(\mathrm{Cay}(\Pi A_n, Z), 1) \to \pi_1(\mathfrak{B}_n^\pi, x_1)$$

be the map of fundamental groups induced by Ψ. Explicitly, the image of a loop $z_1 \dots z_k$ ($z_i \in Z^{\pm 1}$) in $\pi_1(\mathrm{Cay}(\Pi\mathrm{A}_n, Z), 1)$ under Ψ_* is

$$x_1 - z_1(x_1) - z_1 z_2(x_1) - \dots - z_1 z_2 \dots z_k(x_1) = x_1.$$

We first show that Ψ_* is the trivial map, then show that it is also surjective.

Recall that the Cayley graph C of a group G with presentation $\langle X \mid R \rangle$ forms the 1-skeleton of its *Cayley complex*, which we obtain by attaching disks along the loops in C corresponding to all conjugates in G of the words in R. It is well-known that the Cayley complex of a group G is simply-connected [42, Proposition 4.2]. We now verify that the loops in $\pi_1(\mathrm{Cay}(\Pi\mathrm{A}_n, Z), 1)$ corresponding to the relators given at the start of Section 2.4.1 have trivial image under Ψ_*. This allows us to extend Ψ to a map from the (simply-connected) Cayley complex of $\Pi\mathrm{A}_n$ (rel. Z), and so conclude that Ψ_* is trivial.

Note that in the following we confuse a relator with the loop in $\pi_1(\mathrm{Cay}(\Pi\mathrm{A}_n, Z), 1)$ to which it corresponds. Many of the relators $1 - 17$ map to x_1 in $\mathfrak{B}_n^\pi$, as they are words on members of Z that fix x_1. The only ones we need to check are $1 - 3$ and $14 - 17$. Relators $1 - 3$ map into the contractible simplex spanned by $x_1, \dots, x_n$, so are trivial. Relators 14 and 15 are mapped into the simplices $x_1 - x_3$ and $x_1 - x_4$, respectively. We rewrite relators 16 and 17 as $P_{21} w = w P_{21}$ and $\iota_2 w = w \iota_2$. It is clear, then, that relators of type 16 map into the contractible subcomplex of $\mathfrak{B}_n^\pi$ spanned by $x_1, \dots, x_n$ and $x_1 x_2 x_1$, and relators of type 17 map into the contractible subcomplex spanned by $x_1, x_2^{\pm 1}, \dots, x_n$. All relators have now been dealt with, so we conclude that Ψ_* is the trivial map.

We argue as in Day-Putman's proof [24] for the surjectivity of Ψ_*. We represent a loop $\omega \in \pi_1(\mathfrak{B}_n^\pi, x_1)$ as

$$x_1 = w_0 - w_1 - \dots - w_k = x_1,$$

for some $k \geq 0$. We will demonstrate that for any such path (not necessarily with $w_k = x_1$), there exist $\phi_1, \dots, \phi_k \in \Pi\mathrm{A}_n(1)$ such that

$$w_i = \phi_1 t_{12} \phi_2 t_{12} \dots \phi_i t_{12}(x_1),$$

for $0 \leq i \leq k$. We use induction. In the case $k = 0$, there is nothing to prove. Now suppose $k > 0$. Consider the subpath

$$w_0 - w_1 - \dots - w_{k-1}.$$

By induction, to prove the claim all we need find is $\phi_k \in \Pi A_n(1)$ such that

$$w_k = \phi_1 t_{12} \dots \phi_k t_{12}(x_1).$$

We know that $w_{k-1} = \phi_1 t_{12} \dots \phi_{k-1} t_{12}(x_1)$ and w_k form a partial π-basis, therefore so do x_1 and $(\phi_1 t_{12} \dots \phi_{k-1} t_{12})^{-1}(w_k)$. By construction, the action of ΠA_n is transitive on the set of two-element partial π-bases, so there exists $\phi_k \in \Pi A_n(1)$ mapping x_2 to $(\phi_1 t_{12} \dots \phi_{k-1} t_{12})^{-1}(w_k)$. Therefore

$$w_k = \phi_1 t_{12} \dots \phi_k t_{12}(x_1),$$

as required.

Now, we define

$$\phi_{k+1} = (\phi_1 t_{12} \dots \phi_k t_{12})^{-1},$$

so that

$$R := \phi_1 t_{12} \dots \phi_k t_{12} \phi_{k+1} = 1$$

is a relation in ΠA_n. Observe that since $w_k = x_1$, we have $\phi_{k+1} \in \Pi A_n(1)$. Also, the generating set Z contains a subset that generates $\Pi A_n(1)$, by Proposition 2.2.8. We are thus able to write

$$\phi_i = z_1^i \dots z_{p_i}^i,$$

for some $z_j^i \in Z^{\pm 1}$ $(1 \leq i \leq k+1,\ 1 \leq j \leq p_i)$, each of which fixes x_1. We see that $R \in \pi_1(\mathrm{Cay}(\Pi A_n, Z), 1)$ maps to $\omega \in \pi_1(\mathfrak{B}_n^\pi, x_1)$: removing repeated vertices, R maps to

$$x_1 - \phi_1 t_{12}(x_1) - \dots - \phi_1 t_{12} \dots \phi_k t_{12} \phi_{k+1}(x_1) = x_1,$$

which equals ω by construction. Hence Ψ_* is surjective as well as trivial, so $\pi_1(\mathfrak{B}_n^\pi, x_1) = 1$. □

2.4.2 The connectivity of $\mathfrak{B}_n^\pi / \mathcal{PI}_n$

A complex analogous to $\mathfrak{B}_n^\pi$ may be defined when working over $\mathbb{Z}^n$ rather than F_n. We write $\mathcal{B}_n(\mathbb{Z})$ for the *complex of partial bases of* $\mathbb{Z}^n$, whose $(k-1)$-simplices correspond to subsets $\{v_1, \dots, v_k\}$ of free abelian bases of $\mathbb{Z}^n$. Writing members of $\mathbb{Z}^n$ multiplicatively, there is an analogous notion of a palindrome on some fixed free abelian basis V, and so also of a partial π-basis. The *complex of partial π-bases of $\mathbb{Z}^n$* is defined in the obvious way, and denoted $\mathfrak{B}_n^\pi(\mathbb{Z})$.

We first show that $\mathfrak{B}_n^\pi/\mathcal{PI}_n \cong \mathfrak{B}_n^\pi(\mathbb{Z})$, then show that $\mathfrak{B}_n^\pi(\mathbb{Z})$ is $(n-3)$-connected using a related complex studied by Charney [14]. To prove the former, the following lemma is required.

Lemma 2.4.1. *Fix $\{v_1, \ldots, v_n\}$ as a π-basis for $\mathbb{Z}^n$, and let $\rho : F_n \to \mathbb{Z}^n$ be the abelianisation map. Let $\tilde{V} = \{\tilde{v}_1, \ldots, \tilde{v}_k\}$ be a partial π-basis of F_n such that $\rho(\tilde{v}_i) = v_i$ for each $1 \leq i \leq k$. Then we can extend $\tilde{V}$ to a π-basis of F_n, $\{\tilde{v}_1, \ldots, \tilde{v}_n\}$, such that $\rho(\tilde{v}_i) = v_i$ for $1 \leq i \leq n$.*

Proof. Extend $\{\tilde{v}_1, \ldots, \tilde{v}_k\}$ to a full π-basis of F_n, $\{\tilde{v}_1, \ldots, \tilde{v}'_{k+1}, \ldots, \tilde{v}'_n\}$, and define $v'_j = \rho(\tilde{v}'_j)$ for $k+1 \leq j \leq n$. Then $\{v_1, \ldots, v_k, v'_{k+1}, \ldots, v'_n\}$ is a π-basis for $\mathbb{Z}^n$. The group $\Gamma_n[2]$ acts transitively on the set of π-bases of $\mathbb{Z}^n$, so there exists $\phi \in \Gamma_n[2](k)$ such that $\phi(v'_j) = v_j$ for $k+1 \leq j \leq n$. By Proposition 2.3.5, ϕ lifts to some $\tilde{\phi} \in \mathrm{P\Pi A}_n(k)$, and the π-basis $\{\tilde{v}_1, \ldots, \tilde{v}_k, \tilde{\phi}(\tilde{v}'_{k+1}), \ldots, \tilde{\phi}(\tilde{v}'_n)\}$ projects onto $\{v_1, \ldots, v_n\}$ as desired. □

Now we establish an isomorphism of simplicial complexes between $\mathfrak{B}_n^\pi/\mathcal{PI}_n$ and $\mathfrak{B}_n^\pi(\mathbb{Z})$.

Theorem 2.4.2. *The spaces $\mathfrak{B}_n^\pi/\mathcal{PI}_n$ and $\mathfrak{B}_n^\pi(\mathbb{Z})$ are isomorphic as simplicial complexes.*

Proof. Let $\rho : F_n \to \mathbb{Z}^n$ be the abelianisation map, and define a map of simplicial complexes $\Phi : \mathfrak{B}_n^\pi \to \mathfrak{B}_n^\pi(\mathbb{Z})$ on simplices by $\{w_1, \ldots, w_k\} \mapsto \{\rho(w_1), \ldots, \rho(w_k)\}$, for $1 \leq k \leq n$. The map Φ is surjective: by Lemma 2.4.1, each π-basis of $\mathbb{Z}^n$ is projected onto by some π-basis of F_n, and π-bases of $\mathbb{Z}^n$ correspond to maximal simplices of $\mathfrak{B}_n^\pi(\mathbb{Z})$.

It is clear that the map Φ is invariant under the action of $\mathcal{PI}_n$ on $\mathfrak{B}_n^\pi$, and so Φ factors through $\mathfrak{B}_n^\pi/\mathcal{PI}_n$. To establish the theorem, all we need do is show that the induced map from $\mathfrak{B}_n^\pi/\mathcal{PI}_n \to \mathfrak{B}_n^\pi(\mathbb{Z})$ is injective. In other words, we must show that if two simplices s, s' of $\mathfrak{B}_n^\pi$ have the same image under Φ, then s and s' differ by the action of some member of $\mathcal{PI}_n$.

Suppose that $s = \{w_1, \ldots, w_k\}$ and $s' = \{w'_1, \ldots, w'_k\}$ have the same image under Φ. We may assume that $\rho(w_i) = \rho(w'_i)$ for $1 \leq i \leq k$. Let $\Phi(s) = \{\bar{w}_1, \ldots, \bar{w}_k\}$, and extend this partial π-basis of $\mathbb{Z}^n$ to a full π-basis, $W = \{\bar{w}_1, \ldots, \bar{w}_n\}$. By Lemma 2.4.1, we may extend $\{w_1, \ldots, w_k\}$ to $\{w_1, \ldots, w_n\}$ and $\{w'_1, \ldots, w'_k\}$ to $\{w'_1, \ldots, w'_n\}$, such that both of these full π-bases map onto W. Define $\theta \in \Pi\mathrm{A}_n$ by $\theta(w_i) = w'_i$ for $1 \leq i \leq n$. By construction, $\theta(s) = s'$ and $\theta \in \mathcal{PI}_n$, so the theorem is proved. $\square$

This more explicit description of $\mathfrak{B}_n^\pi/\mathcal{PI}_n$ as $\mathfrak{B}_n^\pi(\mathbb{Z})$ enables us to investigate the quotient's connectivity.

Proof of Theorem 2.3.2. By a *unimodular sequence* in $\mathbb{Z}^n$, we mean an (ordered) sequence $(w_1, \ldots, w_k) \subset (\mathbb{Z}^n)^k$ whose entries form a basis of a direct summand of $\mathbb{Z}^n$. Observe that this is just an ordered version of the notion of a partial basis of $\mathbb{Z}^n$. The set of all such sequences of length at least one form a poset under subsequence inclusion. Charney [14] considers (among others) the subposet of sequences $(w_1, \ldots, w_k)$ such that each w_i is congruent to a standard basis vector v_j under mod 2 reduction of the entries of w_i. We denote by $\mathcal{X}_n$ the poset complex given by the subposet of such sequences. Theorem 2.5 of Charney [14] says that $\mathcal{X}_n$ is $(n-3)$-connected.

Let $\mathfrak{B}_n^\pi(\mathbb{Z})^*$ denote the barycentric subdivision of $\mathfrak{B}_n^\pi(\mathbb{Z})$. Label each vertex of $\mathfrak{B}_n^\pi(\mathbb{Z})^*$ by the partial π-basis associated to the simplex of $\mathfrak{B}_n^\pi(\mathbb{Z})$ to which the vertex corresponds. Define a simplicial map $h : \mathcal{X}_n \to \mathfrak{B}_n^\pi(\mathbb{Z})^*$ by $(w_1, \ldots, w_k) \mapsto \{w_1, \ldots, w_k\}$. We may think of h as 'forgetting the order' of each unimodular sequence. Comparing the definitions of $\mathcal{X}_n$ and $\mathfrak{B}_n^\pi(\mathbb{Z})$, it is not immediately clear that h is well-defined: there might be some vertex $(v_1, \ldots, v_k)$ of $\mathcal{X}_n$ such that $\{v_1, \ldots, v_k\}$ extends to a full basis of $\mathbb{Z}^n$, but not a full π-basis. However, viewing the full basis of $\mathbb{Z}^n$ as a matrix in $\Gamma_n[2]$, a straightforward column operations argument shows that this cannot be the case, so h is well-defined.

We see that h induces a map $\pi_i(\mathcal{X}_n) \to \pi_i(\mathfrak{B}_n^\pi(\mathbb{Z})^*)$ for $i \geq 0$, and show that the induced map is surjective. Set a consistent lexicographical order on the vertices of $\mathfrak{B}_n^\pi(\mathbb{Z})^*$, and view $\omega \in \pi_i(\mathfrak{B}_n^\pi(\mathbb{Z})^*)$ as a simplicial i-sphere. The chosen lexicographical ordering allows us to lift ω to $\pi_i(\mathcal{X}_n)$, so the induced maps are surjective. The statement of the theorem follows immediately, as $\pi_i(\mathcal{X}_n) = 1$ for $0 \leq i \leq n-3$. $\square$

2.5 A presentation for $\Gamma_3[2]$

In order to apply Armstrong's theorem [2], it must be the case that $\mathfrak{B}_n^\pi/\mathcal{PI}_n \cong \mathfrak{B}_n^\pi(\mathbb{Z})$ is simply-connected. However, we know from Corollary 2.3.8 that when $n = 3$, the space $\mathfrak{B}_n^\pi(\mathbb{Z})$ has non-trivial fundamental group. The case $n = 3$ forms the base case of our inductive proof of Theorem 2.1.1, so we require an alternative approach to find a generating set for $\mathcal{PI}_3$. Our approach is to find a specific finite presentation of $\Gamma_3[2]$, and use the short exact sequence

$$1 \longrightarrow \mathcal{PI}_3 \longrightarrow \mathrm{P\Pi A}_3 \longrightarrow \Gamma_3[2] \longrightarrow 1$$

to lift the relators in the presentation of $\Gamma_3[2]$ to a normal generating set for $\mathcal{PI}_3$.

2.5.1 A presentation theorem

In order to present $\Gamma_3[2]$, we apply a theorem of Armstrong [3] (not Theorem 2.3.3!) to the action of $\Gamma_3[2]$ on a simply-connected simplicial complex we construct by adding simplices to $\mathfrak{B}_3^\pi(\mathbb{Z})$. Before discussing this complex and the action on it, we introduce the terminology necessary to state and apply Armstrong's theorem. Note that the theorem was first obtained by Brown [12] in greater generality, however Armstrong's theorem is stated more simply, and suffices for our purposes.

Let G be a group acting simplicially on a non-empty, simply-connected simplicial complex K. We shall assume that G does not invert any edges of K. We denote by K^1 the 1-skeleton of K. For a vertex v of K, we denote by G_v its stabiliser subgroup in G. Similarly, we write g_v for any $g \in G$ that stabilises v, to distinguish it as a member of G_v.

We choose a maximal tree M in the graph K^1/G. We lift M to a subtree T in K^1, and take the vertices of T as a set of representatives for the vertices of K/G. Consider an (oriented) edge f of K^1/G not in the tree M. The edge f has a canonical lift to K, e say, such that $i(e) \in T$. Moreover, there is a unique vertex $y \in T$ which is equivalent to $t(e)$ under the action of G. We fix $\gamma_f \in G$ as some member that takes y to $t(e)$. With these choices made, we lift the reverse edge $\bar{f} \in K^1/G$ to $\gamma_f^{-1}(\bar{e})$: this edge has initial vertex $y \in T$ and $\gamma_{\bar{f}} := \gamma_f^{-1}$ sends $i(e)$ to to the terminal vertex of $\gamma_f^{-1}(\bar{e})$. We formally define $\gamma_f = 1$ for $f \in M$.

Let $g \in G$, and let $E = e_1e_2 \dots e_k$ be a sequence of edges joining a fixed vertex $v \in T$ to the vertex $g(v)$. If this path never leaves the tree T, then we must have $g(v) = v$, as the vertices of T are in one-to-one correspondence with the orbits of the action of G on the vertices of K. Suppose then that the path E does not lie completely in T. Then there is a first edge, e_l say, that lies outside T. Armstrong calls the subpath $e_le_{l+1} \dots e_k$ the *tail* of E.

Let $y_1 \in T$ be the initial vertex of e_l. We map e_l into the quotient K^1/G to f_1, say, and consider e^1, the canonical lift of f_1 to K^1. Note that necessarily $i(e^1) = y_1$. Pick some $a_{y_1} \in G_{y_1}$ that maps e^1 to e_l. Such an a_{y_1} must exist, as e^1 and e_l share the initial vertex y_1, and lie in the same orbit under G.

For $l + 1 \le i \le k$, we define

$$e_i^1 = \gamma_{f_1}^{-1} a_{y_1}^{-1}(e_i),$$

and replace the original path E with the newly-constructed path

$$E^1 := e_{l+1}^1 e_{l+2}^1 \dots e_k^1.$$

Armstrong [3] refers to replacing the tail of the path E with a new path as *tail wagging*. Observe that E^1 begins at the vertex $i(e_{l+1}^1) \in T$ and terminates at $\gamma_{f_1}^{-1} a_{y_1}^{-1} g(v)$. Repeatedly applying this process, we will eventually end up with a path that is completely contained inside the tree T, since at each stage of the procedure the tails of the paths strictly decrease in length. This final path will end at the vertex

$$\gamma_{f_q}^{-1} a_{y_q}^{-1} \dots \gamma_{f_1}^{-1} a_{y_1}^{-1} g(v),$$

where the y_j, f_j and a_{y_j} are chosen specifically as above.

By our previous discussion, since this final path lies entirely in T, it must be the case that

$$\gamma_{f_q}^{-1} a_{y_q}^{-1} \dots \gamma_{f_1}^{-1} a_{y_1}^{-1} g = a_v \in G_v.$$

As Armstrong notes [3], this proves that G is generated by the vertex stabilisers G_v $(v \in T)$ and the symbols γ_f, for edges f of K^1/G, since for any $g \in G$, we can now write

$$g = a_{y_1} \gamma_{f_1} \dots a_{y_q} \gamma_{f_q} a_v.$$

Relations in G may arise by traversing the boundary of a 2-simplex Δ in K, starting at a vertex $v \in T$, and then wagging this closed path. Taking $g = 1$ in the above discussion, we get a path that starts and ends at the same vertex: wagging, we obtain an expression r_Δ that is necessarily equal to the identity. We are now able to state Armstrong's Presentation Theorem, with all the above notation assumed.

Theorem 2.5.1 (Armstrong [3])**.** *Let G be a group acting simplicially on a simply-connected simplicial complex K, where G does not invert any edge of K. Fix T as the lift to K of our maximal tree M of K^1/G. Let F be the free group on the symbols λ_f, where f is an (oriented) edge of K^1/G. Then G is presented by taking the free product*

$$(\mathop{\ast}_{w \in T} G_w) * F$$

and adding the relators:

1. *λ_f, if $f \in M$,*
2. *$\lambda_{\bar{f}}\lambda_f$, for all edges f in K^1/G,*
3. *$\lambda_{\bar{f}} g_x \lambda_f (\gamma_{\bar{f}} g \gamma_f)_z^{-1}$, where e is the canonical lift of f to K, and $g \in G$ fixes e,*
4. *r_Δ^λ, where this word is obtain from r_Δ by changing each γ_f in r_Δ to λ_f, and one such word is taken for each G-orbit of 2-simplices.*

We refer to relators of type 3 as *edge relators*, as they arise due to edge stabiliser subgroups of G. We now describe a corollary of Theorem 2.5.1 which gives a simpler presentation when the quotient K/G has a particular structure. It will be this corollary that we apply to present $\Gamma_3[2]$.

Corollary 2.5.2. *Let G, F and K be as in the statement of Theorem 2.5.1. Assume that the image in K/G of each 2-simplex of K has two edges in the maximal tree M. Then G is presented by taking the free product*

$$(\mathop{\ast}_{w \in T} G_w) * F$$

and adding the relators:

1. *$g_x(g_z)^{-1}$, when g stabilises one of the canonical lifts e.*

Proof. Let $f_1f_2f_3$ be the boundary of a 2-simplex Δ in K/G. Without loss of generality, we may assume f_1, $f_2 \in M$ and $f_3 \notin M$. Lift f_i to the canonical $e_i \in T$ $(1 \leq i \leq 3)$. Note that we have $i(e_3)$, $t(e_3) \in T$. We show that the relation $r_\Delta^\lambda = 1$ in this case gives $\lambda_{f_3} = 1$. Tail wagging the path $e_1e_2e_3$, we get the relation

$$\gamma_{f_3}^{-1} a_{i(e_3)}^{-1} = 1,$$

since e_3 is such that $\gamma_{f_3} = a_{i(e_3)} = 1$. Hence we find $r_\Delta^\lambda = \lambda_{f_3}$.

Since all 2-simplices in K/G have two edges in M, we thus kill all the symbols λ_f. We have also arranged it so that all the symbols γ_f are trivial, so the relators in Theorem 2.5.1 may be replaced with those in the statement of the corollary. $\square$

In practice, when presenting G, we find a generating set for the edge stabiliser subgroups, and adjoin one relator for each generator, as in Corollary 2.5.2.

2.5.2 The augmented partial π-basis complex for $\mathbb{Z}_3$

Recall that $\mathcal{B}_n(\mathbb{Z})$ is the partial basis complex of $\mathbb{Z}^n$. We represent its vertices by column vectors $u = \begin{pmatrix} u^{(1)} \\ \vdots \\ u^{(n)} \end{pmatrix}$. For use in the proof of Theorem 2.5.3, we follow Day-Putman [24] and define the *rank of* u to be $|u^{(n)}|$, and denote it by $R(u)$. Let $\mathcal{Y}$ denote the full subcomplex of $\mathcal{B}_3(\mathbb{Z})$ spanned by $\mathfrak{B}_3^\pi(\mathbb{Z})$ and vertices u for which $u^{(1)}$ and $u^{(2)}$ are odd and $u^{(3)}$ is even. We call $\mathcal{Y}$ the *augmented partial π-basis complex for* $\mathbb{Z}_3$. We now demonstrate that $\mathcal{Y}$ is simply-connected.

Theorem 2.5.3. *The complex $\mathcal{Y}$ is simply-connected.*

Proof. By Theorem 2.5 of Charney [14], we know that $\mathfrak{B}_3^\pi(\mathbb{Z})$ is 0-connected, and hence so is $\mathcal{Y}$. To show that $\mathcal{Y}$ is simply-connected, we adapt the proof of Theorem B of Day-Putman [24].

Let u be a vertex of a simplicial complex C. The *link of u in C*, denoted $\mathrm{lk}_C(u)$, is the full subcomplex of C spanned by vertices joined by an edge to u. Let $v_3 \in \mathbb{Z}^3$ be the standard basis vector with third entry 1 and 0s elsewhere. Observe that for any vertex $u \in \mathcal{Y}$, we have $\mathrm{lk}_{\mathcal{Y}}(u) \cong \mathrm{lk}_{\mathcal{Y}}(v_3)$. This is because there is a transitive, simplicial action on the vertices of $\mathcal{Y}$ by the group generated by $\Gamma_3[2]$ and the matrix

$$E = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

This action is transitive because E acts by sending a vertex in $\mathcal{Y} \setminus \mathfrak{B}_3^{\pi}(\mathbb{Z})$ to a vertex of $\mathfrak{B}_3^{\pi}(\mathbb{Z})$, and because $\Gamma_3[2]$ acts transitively on the vertices of $\mathfrak{B}_3^{\pi}(\mathbb{Z})$.

We begin by establishing that $\mathrm{lk}_{\mathcal{Y}}(v_3)$ is connected (and hence, by the above, so is the link of any vertex). By considering what the columns of $M \in \mathrm{GL}(3, \mathbb{Z})$ whose final column is v_3 must look like, we see that a necessary and sufficient condition for $\begin{pmatrix} u^{(1)} \\ u^{(2)} \\ u^{(3)} \end{pmatrix}$ to be a member of $\mathrm{lk}_{\mathcal{Y}}(v_3)$ is that $\begin{pmatrix} u^{(1)} \\ u^{(2)} \end{pmatrix}$ is a vertex of $\mathcal{B}_2(\mathbb{Z})$. The link $\mathrm{lk}_{\mathcal{Y}}(v_3)$ consists of a copy of $\mathcal{B}_2(\mathbb{Z})$ for each $d \in 2\mathbb{Z}$, with two vertices $u, w \in \mathrm{lk}_{\mathcal{Y}}(v_3)$ being joined by an edge if there is an edge between them in *some* copy of $\mathcal{B}_2(\mathbb{Z})$. Hence $\mathrm{lk}_{\mathcal{Y}}(v_3)$ is connected, though note that its fundamental group is an infinite rank free group.

Now, let $\omega \in \pi_1(\mathcal{Y}, v_3)$. We represent ω by the sequence of vertices

$$w_0 - w_1 - \ldots - w_r,$$

where w_i $(1 \leq i \leq r)$ are vertices of $\mathcal{Y}$, and $w_0 = w_k = v_3$. Our goal is to systematically homotope this loop so that the rank of each vertex in the sequence is 0. Such a loop is contained in $\mathrm{lk}_{\mathcal{Y}}(v_3)$, and so may be contracted to the vertex v_3.

Consider a vertex w_i for some $1 < i < r$, with $R(w_i) \neq 0$. Since $\mathrm{lk}_{\mathcal{Y}}(w_i)$ is connected, there is some path

$$w_{i-1} - q_1 - q_2 \ldots - q_s - w_{i+1}$$

in $\mathrm{lk}_{\mathcal{Y}}(w_i)$, as seen in Figure 2.11. Fix attention on some q_j $(1 \leq j \leq s)$. By the Division Algorithm, there exists $a_j, b_j \in \mathbb{Z}$ such that $R(q_j) = a_j \cdot R(w_i) + b_j$ such that $0 \leq b_j < R(w_i)$. As in the proof of Lemma 2.2.10, we wish to ensure that a_j is even, if possible. In all but

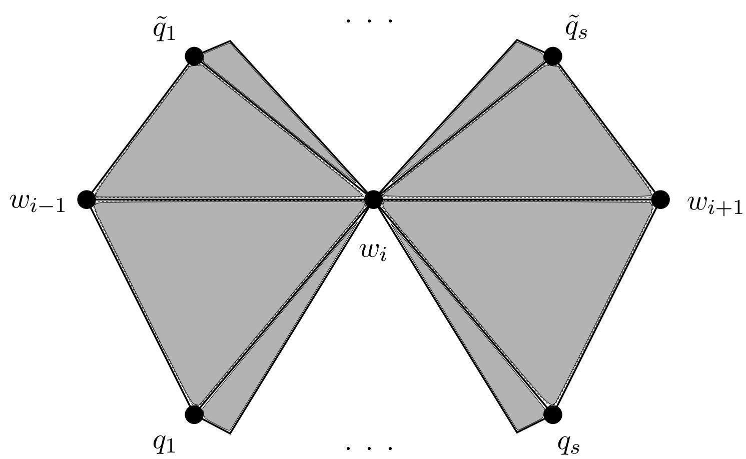

Figure 2.11: *We find two homotopic paths that bound a disk inside* $\mathrm{lk}_{\mathcal{Y}}(w_i)$*, where the 'upper' path seen here is constructed so that* $R(\tilde{q}_j) < R(q_j)$ *for* $1 \leq j \leq s$.

one case, we will be able to rewrite the Division Algorithm as $R(q_j) = A_j \cdot R(w_i) + B_j$, for some $A_j, B_j \in \mathbb{Z}$ such that A_j is even and $0 \leq |B_j| < R(w_i)$. We do a case-by-case parity analysis. Note that $R(q_j)$ and $R(w_i)$ cannot both be odd, as q_j and w_i are joined by an edge. If $R(q_j)$ and $R(w_i)$ have different parities and a_j is odd, we may take $A_j = a_j + 1$ and $B_j = b_j - R(w_i)$. If both $R(q_j)$ and $R(w_i)$ are even, we may still do this, unless $b_j = 0$.

We now associate to each q_j a new vertex, $\tilde{q}_j$, defined by

$$\tilde{q}_j = \begin{cases} q_j - a_j \cdot w_i & \text{if } a_j \text{ even,} \\ q_j - A_j \cdot w_i & \text{if } a_j \text{ odd, } b_j \neq 0, \cdot \\ q_j - a_j \cdot w_i & \text{if } a_j \text{ odd, } b_j = 0 \end{cases}$$

Note that when $b_j = 0$, $R(\tilde{q}_j) = 0$, and under the conditions given, $\tilde{q}_j$ is always well-defined as a vertex of $\mathcal{Y}$. The path

$$w_{i-1} - q_1 - \ldots - q_s - w_{i+1}$$

is homotopic inside $\mathrm{lk}_{\mathcal{Y}}(w_i)$ to the path

$$w_{i-1} - \tilde{q}_1 - \ldots - \tilde{q}_s - w_{i+1},$$

as seen in Figure 2.11. By construction, $R(\tilde{q}_j) < R(w_i)$. Iterating this procedure continually homotopes ω until it is inside $\mathrm{lk}_{\mathcal{Y}}(e_3)$, and hence is trivial. Therefore $\pi_1(\mathcal{Y}) = 1$. □

The complex $\mathfrak{B}_3^{\pi}(\mathbb{Z})$ is not simply-connected. It may be tempting to try to use the method in the above proof to show that $\mathfrak{B}_3^{\pi}(\mathbb{Z})$ is simply-connected, however we know by

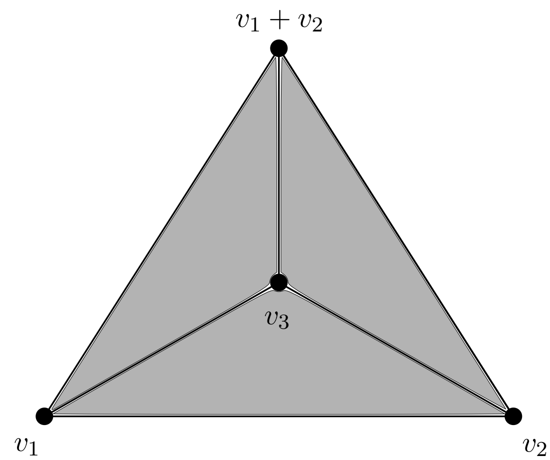

Figure 2.12: *The quotient complex of $\mathcal{Y}$ under the action of $\Gamma_3[2]$. We have labelled its vertices using representatives from the vertex set of $\mathcal{Y}$.*

Corollary 2.3.8 that $\mathfrak{B}_3^\pi(\mathbb{Z})$ has non-trivial fundamental group. The obstruction to the above proof going through occurs when defining $\tilde{q}_j$ in the case that a_j is odd and $b_j = 0$: we find $\tilde{q}_j \notin \mathfrak{B}_3^\pi(\mathbb{Z})$. When a_j is odd and $b_j = 0$, there is no even multiple of w_i that can be added to q_j to decrease its rank, so this method of homotoping loops to a point will not work.

2.5.3 Presenting $\Gamma_3[2]$

Having demonstrated that $\mathcal{Y}$ is simply-connected, we now turn our attention to the obvious action of $\Gamma_3[2]$ on $\mathcal{Y}$. This action is simplicial, does not invert edges, and the quotient complex under the action is contractible, as seen in Figure 2.12. We now apply Corollary 2.5.2 to obtain a presentation of $\Gamma_3[2]$.

We may choose a maximal tree in the 1-skeleton of the quotient of $\mathcal{Y}$ that includes two edges from every 2-simplex, so we are able to use the simpler presentation given by Corollary 2.5.2. We begin by giving a finite presentation for $\Gamma_3[2](v_1)$, the stabiliser of the vertex v_1 in $\mathcal{Y}$, which we obtain from the semi-direct production decomposition of $\Gamma_3[2](1)$ given by Lemma 2.3.5. The group $\Gamma_3[2](v_1)$ is generated by the set $\{O_2, O_3, S_{23}, S_{32}, S_{12}, S_{13}\}$, with a complete list of relators given by

1. ${O_2}^2$,
2. ${O_3}^2$,
3. $(O_2O_3)^2$,
4. $(O_2S_{23})^2$,
5. $(O_2S_{32})^2$,
6. $(O_2S_{12})^2$,
7. $[O_2, S_{13}]$,
8. $(O_3S_{23})^2$,
9. $(O_3S_{32})^2$,
10. $(O_3S_{13})^2$,
11. $[O_3, S_{12}]$,
12. $[S_{23}, S_{13}]$,
13. $[S_{12}, S_{32}]$,
14. $[S_{12}, S_{13}]$,
15. $[S_{12}, S_{23}]S_{13}^{-2}$,
16. $[S_{13}, S_{32}]S_{12}^{-2}$.

By permuting the indices accordingly, we also obtain finite presentations for $\Gamma_3[2](v_2)$ and $\Gamma_3[2](v_3)$. Temporarily ignoring the vertex $v_1 + v_2$ in the quotient, it is clear that the effect of including the edge relators of Corollary 2.5.2 in the presentation of the free product of the stabilisers $\{\Gamma_3[2](v_i)\}$ $(1 \leq i \leq 3)$ produces the presentation given in Corollary 2.1.2 without relators 7 and 10. We denote this (incomplete) presentation by $\mathcal{P}$. This results simply because the stabiliser subgroup of the edge between v_i and v_j $(i \neq j)$ is generated by the intersection of the generating sets for $\Gamma_3[2](v_i)$ and $\Gamma_3[2](v_j)$ given above. For example, the stabiliser $\Gamma_3[2](v_1, v_2)$ of the edge between v_1 and v_2 is generated by $\{O_3, S_{13}, S_{23}\}$.

For the final vertex, $v_1 + v_2$, we begin by abstractly presenting its stabiliser, $\Gamma_3[2](v_1 + v_2)$. Since $\Gamma_3[2](v_1 + v_2)$ and $\Gamma_3[2](v_1)$ are conjugate inside $\mathrm{GL}(3, \mathbb{Z})$, we take the set of formal symbols

$$\{\hat{O}_2, \hat{O}_3, \hat{S}_{23}, \hat{S}_{32}, \hat{S}_{12}, \hat{S}_{13}\}$$

as a generating set for $\Gamma_3[2](v_1 + v_2)$. A defining list of relators for these generators is obtained from the above list for $\Gamma_3[2](v_1)$ by placing a 'hat' above each generator.

The members of $\Gamma_3[2](v_1 + v_2)$ are not, however, strings of formal symbols, but are members of $\Gamma_3[2]$. To express them as such, we observe that

$$\Gamma_3[2](v_1 + v_2) = E_{21} \cdot \Gamma_3[2](v_1) \cdot {E_{21}}^{-1},$$

where E_{21} is the elementary matrix with a 1 in the $(2, 1)$ position. In Table 2.2, we see the conjugates of the generators of $\Gamma_3[2](v_1)$ by E_{21}: these give expressions for the formal symbols generating $\Gamma_3[2](v_1+v_2)$. For example,

$$\hat{S}_{12} = E_{21}S_{12}{E_{21}}^{-1} = O_1O_2S_{21}{S_{12}}^{-1}.$$

Generator M of $\Gamma_3[2](v_1)$	The conjugate $\hat{M} = E_{21} \cdot M \cdot {E_{21}}^{-1}$
O_2	$S_{21}O_2$
O_3	O_3
S_{12}	$O_1O_2S_{21}{S_{12}}^{-1}$
S_{13}	$S_{13}S_{23}$
S_{23}	S_{23}
S_{32}	$S_{32}{S_{31}}^{-1}$

Table 2.2: *The conjugates of the generating set of* $\Gamma_3[2](v_1)$ *by* E_{21}.

We now consider the edge relators corresponding to the final three edges of the quotient of $\mathcal{Y}$. Let f_i be the edge joining v_1+v_2 to v_i ($1 \le i \le 3$), and let J_i be the stabiliser of f_i. We consider these each in turn. Observe that

$$J_2 = E_{21} \cdot \Gamma_3[2](v_1, v_2) \cdot {E_{21}}^{-1},$$

so J_2 is generated by $\{O_3, S_{13}S_{23}, S_{23}\}$. We have expressed those three generators in terms of the generators of $\Gamma_3[2](v_1)$: to obtain our edge relations, we must express them using the generators of $\Gamma_3[2](v_1+v_2)$, and set them to be equal accordingly. Consulting Table 2.2, we get the edge relations $\hat{O}_3 = O_3$, $\hat{S}_{13} = S_{13}S_{23}$ and $\hat{S}_{23} = S_{23}$. Note that these relations simply reiterate the expressions we had already determined for $\hat{O}_3$, $\hat{S}_{13}$ and $\hat{S}_{23}$. Similarly, as we obtain J_3 by conjugating $\Gamma_3[2](v_1, v_3)$ by E_{21}, the edge relations arising from the edge f_3 are $\hat{O}_2 = S_{21}O_2$, $\hat{S}_{12} = O_1O_2S_{21}{S_{12}}^{-1}$ and $\hat{S}_{32} = S_{32}{S_{31}}^{-1}$.

Finally, to obtain J_1, we conjugate $\Gamma_3[2](v_1, v_2)$ by the elementary matrix E_{12}. We obtain that J_1 is generated by $\{O_3, S_{13}, S_{13}S_{23}\}$, which gives edge relations $\hat{O}_3 = O_3$, $S_{13} = \hat{S}_{13}\hat{S}_{23}^{-1}$ and $\hat{S}_{13} = S_{13}S_{23}$. Note that these relations all arise as consequences of the edge relations coming from the edges f_2 and f_3, so are not required.

Thus, using Tietze transformations, we deduce that $\Gamma_3[2]$ is finitely presented by adding the relators in the right-hand column of Table 2.3 to those in the presentation $\mathcal{P}$. Direct

Relator R	Relator R'
$\hat{O}_2^2$	$(S_{21}O_2)^2$
$\hat{O}_3^2$	O_3^2
$(\hat{O}_2\hat{O}_3)^2$	$(S_{21}O_2O_3)^2$
$(\hat{O}_2\hat{S}_{23})^2$	$(S_{21}O_2S_{23})^2$
$(\hat{O}_2\hat{S}_{32})^2$	$(S_{21}O_2S_{32}S_{31}{}^{-1})^2$
$(\hat{O}_2\hat{S}_{12})^2$	$(S_{21}O_2O_1O_2S_{21}S_{12}{}^{-1})^2$
$[\hat{O}_2, \hat{S}_{13}]$	$[S_{21}O_2, S_{13}S_{23}]$
$(\hat{O}_3\hat{S}_{23})^2$	$(O_3S_{23})^2$
$(\hat{O}_3\hat{S}_{32})^2$	$(O_3S_{32}S_{31}{}^{-1})^2$
$(\hat{O}_3\hat{S}_{13})^2$	$(O_3S_{13}S_{23})^2$
$[\hat{O}_3, \hat{S}_{12}]$	$[O_3, O_1O_2S_{21}S_{12}{}^{-1}]$
$[\hat{S}_{23}, \hat{S}_{13}]$	$[S_{23}, S_{13}S_{23}]$
$[\hat{S}_{12}, \hat{S}_{32}]$	$[O_1O_2S_{21}S_{12}{}^{-1}, S_{32}S_{31}{}^{-1}]$
$[\hat{S}_{12}, \hat{S}_{13}]$	$[O_1O_2S_{21}S_{12}{}^{-1}, S_{13}S_{23}]$
$[\hat{S}_{12}, \hat{S}_{23}]\hat{S}_{13}^{-2}$	$[O_1O_2S_{21}S_{12}{}^{-1}, S_{23}](S_{13}S_{23})^{-2}$
$[\hat{S}_{13}, \hat{S}_{32}]\hat{S}_{12}^{-2}$	$[S_{13}S_{23}, S_{32}S_{31}{}^{-1}](O_1O_2S_{21}S_{12}{}^{-1})^{-2}$

Table 2.3: *Here we see two different expressions for each relator: the left-hand column contains expressions in terms of the abstract symbols generating $\Gamma_3[2]$, while the right-hand column reinterprets these relators in terms of the generating set of the presentation $\mathcal{P}$.*

computation reveals that all the relators in the right-hand column, except for the final one, are consequences of the relators of $\mathcal{P}$, so may be removed from our presentation. Let

$$\chi := [S_{13}S_{23}, S_{32}S_{31}{}^{-1}](O_1O_2S_{21}S_{12}{}^{-1})^{-2}$$

be the final relator in Table 2.3. Observe that we may replace the relator in the final entry of the left-hand column with

$$\hat{S}_{32}\hat{S}_{13}\hat{S}_{12}\hat{S}_{32}^{-1}\hat{S}_{13}^{-1}\hat{S}_{12},$$

by rearranging the original relator using relations in $\Gamma_3[2](v_1 + v_2)$. This replaces χ with the relator

$$\chi' := (S_{32}{S_{31}}^{-1}S_{13}S_{23}S_{21}{S_{12}}^{-1})^2,$$

or, more correctly, a word also involving O_1 and O_2 which is clearly equal to χ'. We have thus verified that the presentation given in Corollary 2.1.2 is correct in the $n = 3$ case.

Chapter 3

Outer automorphisms of automorphism groups of right-angled Artin groups

3.1 Overview

Let Γ be a finite simplicial graph, with vertex set V. Let E be the edge set of Γ, which we view as subset of $V \times V$: precisely, $(v, w) \in E$ if and only if the vertices v and w are joined by an edge in Γ. When $(v, w) \in E$, we say v is *adjacent* to w, and vice versa. The graph Γ defines the *right-angled Artin group* A_Γ via the presentation

$$\langle v \in V \mid [v, w] = 1 \text{ if } (v, w) \in E \rangle.$$

The class of right-angled Artin groups contains all finite rank free and free abelian groups, and allows us to interpolate between these two classically well-studied classes of groups.

A centreless group G is *complete* if the natural embedding $\mathrm{Inn}(G) \hookrightarrow \mathrm{Aut}(G)$ is an isomorphism. Dyer-Formanek [27] showed that $\mathrm{Aut}(F_n)$ is complete for F_n a free group of rank $n \geq 2$, giving $\mathrm{Out}(\mathrm{Aut}(F_n)) = 1$. Their approach is algebraic, and relies upon a criterion of Burnside [13], which states that a centreless group G is complete if and only if G is normal in $\mathrm{Aut}(\mathrm{Aut}(G))$, embedding G as $\mathrm{Inn}(G) \hookrightarrow \mathrm{Aut}(G) \hookrightarrow \mathrm{Aut}(\mathrm{Aut}(G))$. Bridson-Vogtmann [10] independently proved the completeness of $\mathrm{Aut}(F_n)$ for $n \geq 3$, using geometric methods, and also showed that $\mathrm{Out}(F_n)$ is complete. Their proof uses the actions of $\mathrm{Aut}(F_n)$ and $\mathrm{Out}(F_n)$ on auter space and outer space, respectively.

Although $\mathrm{Aut}(\mathbb{Z}^n) = \mathrm{GL}(n,\mathbb{Z})$ is not complete, as its centre is $\mathbb{Z}/2$, we observe similar behaviour for free abelian groups: that is, we find that $\mathrm{Out}(\mathrm{GL}(n,\mathbb{Z}))$ has small order for all n. Hua-Reiner [35] explicitly determined $\mathrm{Out}(\mathrm{GL}(n,\mathbb{Z}))$; in particular, we have the following theorem.

Theorem 3.1.1 (Hua-Reiner [35])**.** *Let $A_1 \in \mathrm{Aut}(\mathrm{GL}(n,\mathbb{Z}))$ map $M \in \mathrm{GL}(n,\mathbb{Z})$ to $\det(M) \cdot M$. Let $A_2 \in \mathrm{Aut}(\mathrm{GL}(n,\mathbb{Z}))$ map $M \in \mathrm{GL}(n,\mathbb{Z})$ to $(M^{\mathrm{Tr}})^{-1}$, where M^{Tr} denotes the transpose of M. Then*

- *if $n = 2$, $\mathrm{Out}(\mathrm{GL}(n,\mathbb{Z})) = \mathbb{Z}/2 \times \mathbb{Z}/2$,*
- *if $n \geq 3$ is odd, $\mathrm{Out}(\mathrm{GL}(n,\mathbb{Z})) = \mathbb{Z}/2 = \langle \bar{A}_2 \rangle$,*
- *if $n > 2$ is even, $\mathrm{Out}(\mathrm{GL}(n,\mathbb{Z})) = \mathbb{Z}/2 \times \mathbb{Z}/2 = \langle \bar{A}_1, \bar{A}_2 \rangle$,*

where $\bar{A}_i$ denotes the image of A_i in $\mathrm{Out}(\mathrm{GL}(n,\mathbb{Z}))$ ($i = 1, 2$).

To summarise the results of Dyer-Formanek [27], Bridson-Vogtmann [10] and Hua-Reiner [35], we say that for free or free abelian A_Γ, the orders of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ and $\mathrm{Out}(\mathrm{Out}(A_\Gamma))$ are both uniformly bounded above by 4. The main result of this chapter is that no such uniform upper bounds exist when A_Γ ranges over all right-angled Artin groups.

Theorem 3.1.2. *For any $N \in \mathbb{N}$, there exists a right-angled Artin group A_Γ such that $|\mathrm{Out}(\mathrm{Aut}(A_\Gamma))| > N$.*

We give two proofs of Theorem 3.1.2: in the first, we work over right-angled Artin groups with non-trivial centre, while in the second we work over right-angled Artin groups with trivial centre. We also prove the analogous result regarding the order of $\mathrm{Out}(\mathrm{Out}(A_\Gamma))$.

Theorem 3.1.3. *For any $N \in \mathbb{N}$, there exists a right-angled Artin group A_Γ such that $|\mathrm{Out}(\mathrm{Out}(A_\Gamma))| > N$.*

We remark that neither Theorem 3.1.2 nor 3.1.3 follows from the other, since in general, given a quotient G/N, the groups $\mathrm{Aut}(G/N)$ and $\mathrm{Aut}(G)$ may behave very differently.

Many of the groups that arise in geometric group theory display 'algebraic rigidity', in the sense that their outer automorphism groups are small. The aforementioned results of Dyer-Formanek [27], Bridson-Vogtmann [10] and Hua-Reiner [35] are examples of this phenomenon. Further examples are given by braid groups [28] and many mapping class groups [36], as these groups have $\mathbb{Z}/2$ as their outer automorphism groups. Theorems 3.1.2 and 3.1.3 thus fit into a more general framework of the study of algebraic rigidity within geometric group theory.

We prove both theorems by exhibiting classes of right-angled Artin groups over which the groups in question grow without bound. We introduce the notions of an *austere graph* and an *austere graph with star cuts* in Sections 3.2 and 3.4, respectively. These lead to tractable decompositions of $\mathrm{Aut}(A_\Gamma)$ and $\mathrm{Out}(A_\Gamma)$ as semi-direct products, which then yield numerous members of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ and $\mathrm{Out}(\mathrm{Out}(A_\Gamma))$. Our methods do not obviously yield infinite order elements of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$; we discuss this further in Section 3.5.

3.1.1 Outline of chapter

In Section 3.2, we recall a finite generating set of $\mathrm{Aut}(A_\Gamma)$ and give the proof of Theorem 3.1.3. Sections 3.3 and 3.4 contain two proofs of Theorem 3.1.2; first for right-angled Artin groups with non-trivial centre, then for those with trivial centre. In Section 3.5, we discuss generalisations of this work, including the question of extremal behaviour of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$.

3.2 Proof of Theorem 3.1.3

Let Γ be a finite simplicial graph with vertex set V and edge set $E \subset V \times V$. We write $\Gamma = (V, E)$. We will abuse notation and consider $v \in V$ as both a vertex of Γ and a generator of A_Γ. We will also often consider a subset $S \subseteq V$ as the full subgraph of Γ which it spans. For a vertex $v \in V$, we define its *link*, $\mathrm{lk}(v)$, to be the set of vertices in V adjacent to v, and its *star*, $\mathrm{st}(v)$, to be $\mathrm{lk}(v) \cup \{v\}$.

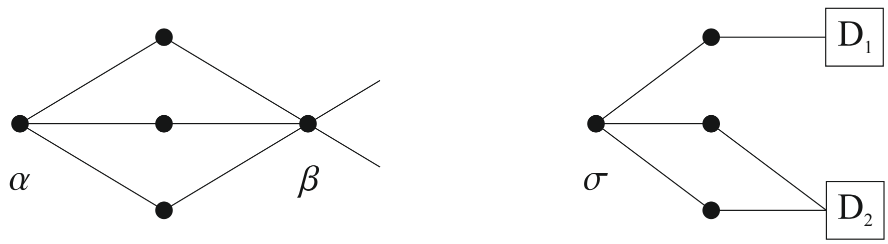

(a) (b)

Figure 3.1: **(a)** *The local picture of a vertex α being dominated by a vertex β.* **(b)** *Removing the star of the vertex σ leaves two connected components,* D_1 *and* D_2.

3.2.1 The LS generators

Laurence [41], proving a conjecture of Servatius [54], gave a finite generating set for $\mathrm{Aut}(A_\Gamma)$, which we now recall. We specify the action of the generator on the elements of V. If a vertex $v \in V$ is omitted, it is assumed to be fixed. There are four types of generators:

1. *Inversions*, ι_v: for each $v \in V$, ι_v maps v to v^{-1}. We denote by I_Γ the subgroup of $\mathrm{Aut}(A_\Gamma)$ generated by the inversions.

2. *Graph symmetries*, ϕ: each $\phi \in \mathrm{Aut}(\Gamma)$ induces an automorphism of A_Γ, which we also denote by ϕ, mapping $v \in V$ to $\phi(v)$.

3. *Dominated transvections*, τ_{xy}: for $x, y \in V$, whenever $\mathrm{lk}(y) \subseteq \mathrm{st}(x)$, we write $y \leq x$, and say y is *dominated* by x (see Figure 3.1a). In this case, τ_{xy} is well-defined, and maps y to yx. The vertex x may be adjacent to y, but it need not be.

4. *Partial conjugations*, $\gamma_{c,D}$: fix $c \in V$, and select a connected component D of $\Gamma \setminus \mathrm{st}(c)$ (see Figure 3.1b). The partial conjugation $\gamma_{c,D}$ maps every $d \in D$ to cdc^{-1}. We denote by $\mathrm{PC}(A_\Gamma)$ the subgroup of $\mathrm{Aut}(A_\Gamma)$ generated by the partial conjugations.

We refer to the generators on this list as the *LS generators* of $\mathrm{Aut}(A_\Gamma)$.

3.2.2 Austere graphs.

We say that a graph $\Gamma = (V, E)$ is *austere* if it has trivial symmetry group, no dominated vertices, and for each $v \in V$, the graph $\Gamma \setminus \mathrm{st}(v)$ is connected. We describe such graphs as being 'austere' because their properties make the list of LS generators of $\mathrm{Aut}(A_\Gamma)$ as short as possible. We use examples of austere graphs to prove Theorem 3.1.3.

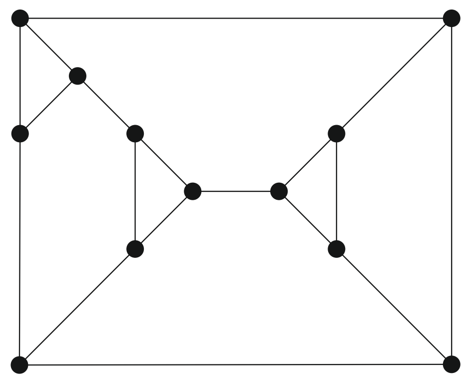

Figure 3.2: *The Frucht graph, an example of a graph which is austere.*

Proof of Theorem 3.1.3. For an austere graph $\Gamma = (V, E)$, the only well-defined LS generators of $\mathrm{Aut}(A_\Gamma)$ are the inversions and the partial conjugations, as by definition Γ has no symmetries or dominated vertices. Let $n = |V|$. Note that each partial conjugation is an inner automorphism, since by construction, for each $v \in V$, the graph $\Gamma \setminus \mathrm{st}(v)$ is connected. We have the decomposition

$$\mathrm{Aut}(A_\Gamma) \cong \mathrm{Inn}(A_\Gamma) \rtimes I_\Gamma,$$

where $I_\Gamma \cong (\mathbb{Z}/2)^n$ is the group generated by the inversions. The inversions act on $\mathrm{Inn}(A_\Gamma) \cong A_\Gamma$ in the obvious way, either inverting or fixing (conjugation by) each $v \in V$. We have

$$\mathrm{Out}(A_\Gamma) = \mathrm{Aut}(A_\Gamma)/\mathrm{Inn}(A_\Gamma) \cong I_\Gamma,$$

and so, since I_Γ is abelian,

$$\mathrm{Aut}(\mathrm{Out}(A_\Gamma)) \cong \mathrm{Out}(\mathrm{Out}(A_\Gamma)) \cong \mathrm{GL}(n, \mathbb{Z}/2).$$

If we can find austere graphs for which n is as large as we like, then we will have proved Theorem 3.1.3, as the order of $\mathrm{GL}(n, \mathbb{Z}/2)$ strictly increases with n.

The Frucht graph, seen in Figure 3.2, was constructed by Frucht [30] as an example of a 3-regular simplicial graph with trivial symmetry group. In fact, it is easily checked that the Frucht graph is austere. Baron-Imrich [4] generalised the Frucht graph to produce a family of finite, 3-regular simplicial graphs with trivial symmetry groups, over which $n = |V|$ is unbounded. Like the Frucht graph, these graphs may also be shown to be austere, and so they define a class of right-angled Artin groups which proves Theorem 3.1.3. □

3.3 Proof of Theorem 3.1.2: right-angled Artin groups with non-trivial centre

In this section, we assume that A_Γ has non-trivial centre. Let $\{\Gamma_i\}$ be a collection of graphs. The *join*, $\mathcal{J}\{\Gamma_i\}$, of $\{\Gamma_i\}$ is the graph obtained from the disjoint union of $\{\Gamma_i\}$ by adding an edge (v_i, v_j) for all vertices v_i of Γ_i and v_j of Γ_j, for all $i \neq j$. Observe that for a finite collection of finite simplicial graphs $\{\Gamma_i\}$, we have

$$A_{\mathcal{J}\{\Gamma_i\}} \cong \prod_i A_{\Gamma_i},$$

as the edges we add to form the join correspond precisely with the relators needed to form the direct product. When we take the join of only two graphs, Γ and Δ, we write $\mathcal{J}(\Gamma, \Delta)$ for their join.

3.3.1 Decomposing $\mathrm{Aut}(A_\Gamma)$

A vertex $s \in V$ is said to be *social* if it is adjacent to every vertex of $V \setminus \{s\}$. Let S denote the set of social vertices of Γ and set $k = |S|$. Let $\Delta = \Gamma \setminus S$. We have $\Gamma = \mathcal{J}(S, \Delta)$, so $A_\Gamma \cong \mathbb{Z}^k \times A_\Delta$, and by *The Centralizer Theorem* of Servatius [54], the centre of A_Γ is $A_S = \mathbb{Z}^k$. A first step to understanding how the structure of $\mathrm{Aut}(A_\Gamma)$ relates to $\mathrm{Aut}(A_S)$ and $\mathrm{Aut}(A_\Delta)$ is the following proposition.

Proposition 3.3.1. *The group* $\mathrm{GL}(k, \mathbb{Z}) \times \mathrm{Aut}(A_\Delta)$ *is a proper subgroup of* $\mathrm{Aut}(A_\Gamma)$.

Proof. We examine the LS generators that are well-defined for each graph, S and Δ. Each of these LS generators is also a well-defined LS generator for $\mathrm{Aut}(A_\Gamma)$, as we now show by considering each in turn. The inversions are clearly well-defined in $\mathrm{Aut}(A_\Gamma)$. Any $\phi \in \mathrm{Aut}(\Gamma)$ must preserve S and Δ as sets, since clearly ϕ must preserve the property of being (or not being) adjacent to every other vertex of Γ, so all the graph symmetries of S and Δ extend to symmetries of Γ. There are no partial conjugations to consider in $\mathrm{Aut}(A_S)$, and removing the star of any vertex in $\Gamma \setminus S$ produces the same graph as removing the same vertex from Δ (considered as an abstract graph, not a subgraph of Γ), so any partial conjugation of A_Δ extends to a partial conjugation of A_Γ. Let x dominate y in S (resp. Δ). In Γ, x and y both gain all vertices in Δ (resp. S) as neighbours, so x continues to dominate y in Γ. Thus, the dominated transvection τ_{xy} is well-defined in $\mathrm{Aut}(A_\Gamma)$.

It is easy to see that these LS generators generate $\mathrm{GL}(k, \mathbb{Z})$ and $\mathrm{Aut}(A_\Delta)$ inside $\mathrm{Aut}(A_\Gamma)$. That they generate the direct product in the statement of the proposition follows from observing that they act on different sets of vertices: namely, S and Δ. We get a proper subgroup of $\mathrm{Aut}(A_\Gamma)$, as there exist LS generators of $\mathrm{Aut}(A_\Gamma)$ which do not preserve A_S and A_Δ as sets. □

No vertex $v \in \Delta$ can dominate any vertex of S (otherwise v would be social), however each $s \in S$ dominates each $v \in \Delta$: this is due to the join construction of Γ, and also since S consists of social vertices. The only LS generators not contained in the proper subgroup $\mathrm{GL}(k, \mathbb{Z}) \times \mathrm{Aut}(A_\Delta)$ are of the form τ_{sa}, where $s \in S$ and $a \in \Delta$. We will refer to this type of transvection as a *lateral transvection*, as they occur 'between' the two graphs, S and Δ.

Proposition 3.3.2. *Let* $\Gamma = \mathcal{J}(S, \Delta)$ *define a right-angled Artin group,* A_Γ*, with non-trivial centre. The group* $\mathcal{L}$ *generated by the lateral transvections is isomorphic to* $\mathbb{Z}^{k|\Delta|}$.

Proof. It is clear the lateral transvections τ_{sa} and τ_{tb} commute if $a \neq b$, as they act on distinct vertices. The only case left to check is τ_{sa} and τ_{ta}, for $s, t \in S$ and $a \in \Delta$. We see that

$$\tau_{ta}\tau_{sa}\tau_{ta}^{-1}(a) = \tau_{ta}\tau_{sa}(at^{-1}) = \tau_{ta}(ast^{-1}) = atst^{-1} = as,$$

since s and t commute. Therefore $\tau_{ta}\tau_{sa}\tau_{ta}^{-1} = \tau_{sa}$, and hence $\mathcal{L}$ is abelian.

We see that $\mathcal{L}$ is torsion-free: suppose $T \in \mathcal{L}$ sends $a \in \Delta$ to aw, for some $w \in A_S = \mathbb{Z}^k$. Let $m \in \mathbb{Z}$ such that T^m is the identity. Then $T^m(a) = aw^m = a$. Since $\mathbb{Z}^k$ is torsion-free, we must have $m = 0$, and so $\mathcal{L}$ is torsion-free.

A straightforward calculation verifies that the lateral transvections form a $\mathbb{Z}$-basis for $\mathcal{L}$: suppose $T_1^{p_1} \dots T_r^{p_r} = 1$, where the T_i are lateral transvections and $p_i \in \mathbb{Z}$ $(1 \leq i \leq r)$. As soon as some p_i is non-zero, some vertex $a \in \Delta$ is not fixed, and so the product $T_1^{p_1} \dots T_r^{p_r}$ is non-trivial. This relies upon the fact that S forms a $\mathbb{Z}$-basis for A_S.

To deduce the rank, observe there is a bijection between $\{\tau_{sa} \mid S \in S, a \in \Delta\}$ and $S \times \Delta$. □

We now show that $\mathcal{L}$ is the kernel of a split product decomposition of $\mathrm{Aut}(A_\Gamma)$. This is an $\mathrm{Aut}(A_\Gamma)$ version of a decomposition of $\mathrm{Out}(A_\Gamma)$ given by Charney-Vogtmann [17].

Proposition 3.3.3. *Let $\Gamma = \mathcal{J}(S, \Delta)$ define a right-angled Artin group, A_Γ, with non-trivial centre. The group $\mathrm{Aut}(A_\Gamma)$ splits as the product*

$$\mathbb{Z}^{k|\Delta|} \rtimes [\mathrm{GL}(k, \mathbb{Z}) \times \mathrm{Aut}(A_\Delta)] .$$

Proof. Standard computations show that $\mathcal{L} \cong \mathbb{Z}^{k|\Delta|}$ is closed under conjugation by the LS generators and their inverses: these calculations are summarised in Table 3.1. Note that we decompose any $\phi \in \mathrm{Aut}(\Gamma)$ into its actions on S and Δ, and use a proper subset T of the LS generators, which suffices to generate $\mathrm{Aut}(A_\Gamma)$.

$\lambda \in T \cup T^{-1}$	$\lambda \cdot \tau_{sa} \cdot \lambda^{-1}$	$\lambda \in T \cup T^{-1}$	$\lambda \cdot \tau_{sa} \cdot \lambda^{-1}$
ι_t	τ_{sa}	ι_b	τ_{sa}
ι_s	$-\tau_{sa}$	ι_a	$-\tau_{sa}$
τ_{st}	τ_{sa}	τ_{bd}	τ_{sa}
τ_{rt}	τ_{sa}	τ_{ab}	$\tau_{sa} - \tau_{sb}$
τ_{ts}	$\tau_{sa} + \tau_{ta}$	τ_{ab}^{-1}	$\tau_{sa} + \tau_{sb}$
τ_{ts}^{-1}	$\tau_{sa} - \tau_{ta}$	$\phi \in \mathrm{Aut}(\Delta)$	$\tau_{s\phi(a)}$
		$\gamma_{c,D}$	τ_{sa}

Table 3.1: *The conjugates of a lateral transvection τ_{sa}. The vertices $a, b, d \in \Delta$ and $r, s, t \in S$ are taken to be distinct, and D being any connected component of $\Gamma \setminus \mathrm{st}(c)$.*

We observe that the intersection of $\mathcal{L}$ and $\mathrm{GL}(k, \mathbb{Z}) \times \mathrm{Aut}(A_\Delta)$ is trivial: the elements of $\mathcal{L}$ transvect vertices of Δ by vertices of S, whereas the elements of $\mathrm{GL}(k, \mathbb{Z}) \times \mathrm{Aut}(A_\Delta)$ carry $\mathbb{Z}^k$ and A_Δ back into themselves. Thus, $\mathrm{Aut}(A_\Gamma)$ splits as in the statement of the proposition. □

We look to the $\mathbb{Z}^{k|\Delta|}$ kernel as a source of automorphisms of $\mathrm{Aut}(A_\Gamma)$. We must however ensure that the split product action is preserved; this is achieved using the theory of automorphisms of split products, which we now recall.

3.3.2 Automorphisms of split products

We refer the reader to Brown [11], Passi-Singh-Yadav [50] and Wells [59] for further details of the exposition in this subsection.

Let $G = N \rtimes H$ be a split product, where N is abelian, with the action of H on N being encoded by a homomorphism $\alpha : H \to \mathrm{Aut}(N)$, writing $\alpha(h) = \alpha_h$. For brevity, we will often write $(n,h) \in G$ simply as nh, when the meaning is clear. Let $\mathrm{Aut}(G,N) \leq \mathrm{Aut}(G)$ be the subgroup of automorphisms that preserve N as a set. Let $\gamma \in \mathrm{Aut}(G,N)$. We get an induced automorphism ϕ, say, of G/N, defined by $\phi(gN) = \gamma(g)N$. This is well-defined since $\gamma(N) = N$. We also obtain an automorphism θ, say, of N, by restriction: that is, $\theta = \gamma|_N$. The map $P : \mathrm{Aut}(G,N) \to \mathrm{Aut}(N) \times \mathrm{Aut}(H)$ given by $P(\gamma) = (\theta, \phi)$ is a homomorphism.

An element $(\theta,\phi) \in \mathrm{Aut}(N) \times \mathrm{Aut}(H)$ is said to be a *compatible pair* if $\theta\alpha_h\theta^{-1} = \alpha_{\phi(h)}$, for all $h \in H$. Let $C \leq \mathrm{Aut}(N) \times \mathrm{Aut}(H)$ be the subgroup of all compatible pairs. This is a special (split, abelian kernel) case of the notion of compatibility for group extensions [50], [59]. Notice that the image of P is contained in C, since $\gamma \in \mathrm{Aut}(G,N)$ must preserve the relation $hnh^{-1} = \alpha_h(n)$ for all $h \in H, n \in N$. We therefore restrict the codomain of P to C. Note that while P (with its new codomain) is surjective, it need not be injective. Injectivity may fail since the map P does not see the difference between automorphisms of G that preserve H and those which do not. Precisely, there may be some $\gamma \in \mathrm{Aut}(G,N)$ which restricts to the identity on N and induces the identity on G/N, so is in the kernel of P, but maps $(1,h)$ to (n_h,h), where n_h need not be trivial.

We map C back into $\mathrm{Aut}(G,N)$ using the homomorphism R, defined by

$$R(\theta,\phi)(nh) = \theta(n)\phi(h).$$

Let $\mathrm{Aut}_H(G,N)$ be the subgroup of $\mathrm{Aut}(G,N)$ of maps which induce the identity on H.

This group is mapped via P onto

$$C_1 := \{\theta \in \mathrm{Aut}(N) \mid \theta\alpha(h)\theta^{-1} = \alpha(h) \ \ \forall h \in H\}.$$

Note C_1 is the centraliser of $\mathrm{im}(\alpha)$ in $\mathrm{Aut}(N)$. Our strategy for proving Theorem 3.1.2 is to determine C_1 for the split decomposition of $\mathrm{Aut}(A_\Gamma)$ given by Proposition 3.3.3, and use R to map C_1 into $\mathrm{Aut}(\mathrm{Aut}(A_\Gamma))$.

3.3.3 Ordering the lateral transvections

In order to determine the image of α for our split product, $\mathbb{Z}^{k|\Delta|} \rtimes [\mathrm{GL}(k,\mathbb{Z}) \times \mathrm{Aut}(A_\Delta)]$, we specify an ordering on the lateral transvections. We do this because the image of α is a subgroup of $\mathrm{Aut}(\mathbb{Z}^{k|\Delta|}) = \mathrm{GL}(k|\Delta|,\mathbb{Z})$, and the ordering we specify will allow us to give a concrete description of the members of this subgroup using block matrices.

Let $s_1 \leq \ldots \leq s_k$ be an arbitrary total order on the vertices of S. For lateral transvections $\tau_{s_i a}, \tau_{s_j b}$, we say $\tau_{s_i a} \leq \tau_{s_j b}$ if $s_i \leq s_j$. For a fixed i, we refer to the set $\{\tau_{s_i a} \mid a \in \Delta\}$ as a Δ-*block*.

We now use properties of the graph Δ to determine the rest of the ordering on the lateral transvections. Recall that for vertices $x, y \in V$, x dominates y if $\mathrm{lk}(y) \subseteq \mathrm{st}(x)$, and we write $y \leq x$. Charney-Vogtmann [17, Lemma 2.2] show that $\leq$ is a pre-order (that is, a reflexive, transitive relation) on V, and use it to define the following equivalence relation. Let $v, w \in V$. We say v and w are *domination equivalent* if $v \leq w$ and $w \leq v$. If this is the case, we write $v \sim w$, and let $[v]$ denote the domination equivalence class of v.

The pre-order on V descends to a partial order on $V/\sim$, the set of domination equivalence classes of V. [17]. We also denote this partial order by $\leq$. The group $\mathrm{Aut}(\Delta)$ acts on the set of domination classes of Δ. Let $\mathcal{O}$ be the set of orbits of this action, writing $\mathcal{O}_{[v]}$ for the orbit of the class $[v]$. Note that, by construction, there is a transitive action of $\mathrm{Aut}(\Delta)$ on $\mathcal{O}_{[v]}$ for each $v \in \Delta$. We wish to define a partial order $\ll$ on $\mathcal{O}$ which respects the partial order on the domination classes. That is, if $[v] \leq [w]$, then $\mathcal{O}_{[v]} \ll \mathcal{O}_{[w]}$, for domination classes $[v]$ and $[w]$.

We achieve this by defining a relation $\ll$ on $\mathcal{O}$ by the rule $\mathcal{O}_{[v]} \ll \mathcal{O}_{[w]}$ if and only if there exists $[w'] \in \mathcal{O}_{[w]}$ such that $[v] \leq [w']$. This is well-defined, since $\mathrm{Aut}(\Delta)$ acts transitively on each $\mathcal{O}_{[v]} \in \mathcal{O}$. The properties of $\leq$ discussed above give us the following proposition.

Proposition 3.3.4. *The relation* $\ll$ *on* $\mathcal{O}$ *is a partial order.*

Proof. The relation $\ll$ is reflexive, since $[v] \leq [v]$ for all $v \in \Delta$. To obtain transitivity and anti-symmetry of $\ll$, we utilise the transitive action of $\mathrm{Aut}(\Delta)$ on each $\mathcal{O}_{[v]} \in \mathcal{O}$. Suppose $\mathcal{O}_{[u]} \ll \mathcal{O}_{[v]} \ll \mathcal{O}_{[w]}$. By acting on $\mathcal{O}_{[v]}$ and $\mathcal{O}_{[w]}$ by a member of $\mathrm{Aut}(\Delta)$ if need be, we may conclude that $[u] \leq [v] \leq [w]$, so $[u] \leq [w]$. Thus, $\mathcal{O}_{[u]} \ll \mathcal{O}_{[w]}$, and so $\ll$ is transitive. Anti-symmetry of $\ll$ may be established by noting that if $[v] \leq [w]$, then $|\mathrm{st}(v)| \leq |\mathrm{st}(w)|$, and if $[v] \leq [w]$ with $|\mathrm{st}(v)| = |\mathrm{st}(w)|$ then $[v] = [w]$. $\square$

We use $\ll$ to define a total order on the vertices of Δ, by first extending $\ll$ to a total order on $\mathcal{O}$. We also place total orders on the domination classes within each $\mathcal{O}_{[v]} \in \mathcal{O}$, and on the vertices within each domination class. Now each vertex is relabelled $T(p, q, r)$ to indicate its place in the order: $T(p, q, r)$ is the rth vertex of the qth domination class of the pth orbit. Precisely, we order the set of symbols $\{T(p, q, r)\}$ using a lexicographic ordering on the set $\{(p, q, r)\}$. When working with a given Δ-block, we can identify the lateral transvections with the vertices of Δ, allowing us to think of $T(p, q, r)$ as a lateral transvection. Thus, we may think of a specific Δ-block as inheriting an order from the ordering on Δ.

3.3.4 The centraliser of the image of α

We now explicitly determine the image of α, and its centraliser, in $\mathrm{GL}(k|\Delta|, \mathbb{Z})$. Looking at how $\mathrm{GL}(k, \mathbb{Z}) \times \mathrm{Aut}(A_\Delta)$ acts on $\mathbb{Z}^{k|\Delta|}$ (see Table 3.1), we see that the image of α is

$$Q := \mathrm{GL}(k, \mathbb{Z}) \times \Phi_\Delta,$$

where $\Phi_\Delta \leq \mathrm{GL}(|\Delta|, \mathbb{Z})$ is the image of $\mathrm{Aut}(A_\Delta)$ under the homomorphism induced by abelianising A_Δ. The action of Q on $\mathbb{Z}^{k|\Delta|}$ factors through $\mathrm{GL}(k+|\Delta|, \mathbb{Z})$ in an obvious way, as pointed out by an anonymous referee, simply by mapping Q and $\mathbb{Z}^{k|\Delta|}$ into $\mathrm{GL}(k+|\Delta|, \mathbb{Z})$ via this induced homomorphism. Working in $\mathrm{GL}(k + |\Delta|, \mathbb{Z})$ instead of $\mathrm{GL}(k|\Delta|, \mathbb{Z})$ is simpler, however it does not allow us to fully determine the group C_1: working in $\mathrm{GL}(k|\Delta|, \mathbb{Z})$, we exhaust the members of C_1.

The matrices in Q have a natural block decomposition given by the Δ-blocks: each $M \in Q$ may be partitioned into k horizontal blocks and k vertical blocks, each of which has size $|\Delta| \times |\Delta|$. We write $M = (A_{ij})$, where A_{ij} is the block matrix entry in the ith row and jth column. Under this decomposition, we see that the $\mathrm{GL}(k, \mathbb{Z})$ factor of Q is embedded as

$$\mathrm{GL}(k, \mathbb{Z}) \cong \{(a_{ij} \cdot I_{|\Delta|}) \mid (a_{ij}) \in \mathrm{GL}(k, \mathbb{Z})\},$$

where $I_{|\Delta|}$ is the identity matrix in $\mathrm{GL}(|\Delta|, \mathbb{Z})$. We write $\mathrm{Diag}(D_1, \dots, D_k)$ to denote the block diagonal matrix (B_{ij}) where $B_{ii} = D_i$ and $B_{ij} = 0$ if $i \neq j$. The Φ_Δ factor of Q embeds as

$$\Phi_\Delta \cong \{\mathrm{Diag}(M, \dots, M) \mid M \in \Phi_\Delta\} \leq Q.$$

We now determine the centraliser, $C(Q)$, of Q in $\mathrm{GL}(k|\Delta|, \mathbb{Z})$. The proof is similar to the standard computation of $Z(\mathrm{GL}(k, \mathbb{Z}))$.

Lemma 3.3.5. *The centraliser $C(Q)$ is a subgroup of* $\{\mathrm{Diag}(M, \dots, M) \mid M \in \mathrm{GL}(|\Delta|, \mathbb{Z})\}$.

Proof. Clearly an element of $C(Q)$ must centralise the $\mathrm{GL}(k, \mathbb{Z})$ factor of Q. Let D be the subgroup of diagonal matrices in $\mathrm{GL}(k, \mathbb{Z})$, and define

$$\hat{D} := \{(\epsilon_{ij} \cdot I_{|\Delta|}) \mid (\epsilon_{ij}) \in D\} \leq Q.$$

Suppose $(A_{ij}) \in C(Q)$ centralises $\hat{D}$. Then for each $(\epsilon_{ij} \cdot I_{|\Delta|}) \in \hat{D}$, we must have

$$(A_{ij}) = (\epsilon_{ij} \cdot I_{|\Delta|})(A_{ij})(\epsilon_{ij} \cdot I_{|\Delta|}) = (\epsilon_{ii}\epsilon_{jj}A_{ij}),$$

since $(\epsilon_{ij} \cdot I_{|\Delta|})$ is block diagonal. Since $\epsilon_{ii} \in \{-1, 1\}$ for $1 \leq i \leq k$, we must have $A_{ij} = 0$ if $i \neq j$, so (A_{ij}) is block diagonal. By considering which block diagonal matrices centralise $(E_{ij} \cdot I_{|\Delta|})$, where $(E_{ij}) \in \mathrm{GL}(k, \mathbb{Z})$ is an elementary matrix, we see that any block diagonal matrix centralising the $\mathrm{GL}(k, \mathbb{Z})$ factor of Q must have the *same* matrix $M \in \mathrm{GL}(|\Delta|, \mathbb{Z})$ in each diagonal block. It is then a standard calculation to verify that any choice of $M \in \mathrm{GL}(|\Delta|, \mathbb{Z})$ will centralise the $\mathrm{GL}(k, \mathbb{Z})$ factor of Q. □

The problem of determining $C(Q)$ has therefore been reduced to determining the centraliser of Φ_Δ in $\mathrm{GL}(|\Delta|, \mathbb{Z})$. The total order we specified on the vertices of Δ gives a block lower triangular decomposition of $M \in \Phi_\Delta$, which we utilise in the proof of Proposition 3.3.6. This builds upon a matrix decomposition given by Day [25] and Wade [57].

Observe that Φ_Δ contains the diagonal matrices of $\mathrm{GL}(|\Delta|, \mathbb{Z})$. As in the above proof, anything centralising Φ_Δ must be a diagonal matrix. For a diagonal matrix $E \in \mathrm{GL}(|\Delta|, \mathbb{Z})$, we write $E(p, q, r)$ for the diagonal entry corresponding to the vertex $T(p, q, r)$ of Δ.

Proposition 3.3.6. *A diagonal matrix* $E \in \mathrm{GL}(|\Delta|, \mathbb{Z})$ *centralises* Φ_Δ *if and only if the following conditions hold:*

(1) *If* $p = p'$*, then* $E(p, q, r) = E(p', q', r')$*, and,*

(2) *If* $T(p, q, r)$ *is dominated by* $T(p', q', r')$*, then* $E(p, q, r) = E(p', q', r')$

Proof. We define a block decomposition of the matrices in $\mathrm{GL}(|\Delta|, \mathbb{Z})$ using the sizes of the orbits, $\mathcal{O}_{[v_1]} \ll \ldots \ll \mathcal{O}_{[v_l]}$. Let $m_i = |\mathcal{O}_{[v_i]}|$. We partition $M \in \mathrm{GL}(|\Delta|, \mathbb{Z})$ into l horizontal blocks and l vertical blocks, writing $M = (M_{ij})$, where M_{ij} is an $m_i \times m_j$ matrix. Observe that due to the ordering on the lateral transvections, if $i < j$, then $M_{ij} = 0$.

Let $E \in \mathrm{GL}(|\Delta|, \mathbb{Z})$ satisfy the conditions in the statement of the proposition. We may write $E = \mathrm{Diag}(\epsilon_1 \cdot I_{m_1 \times m_1}, \ldots, \epsilon_l \cdot I_{m_l \times m_l})$, where each $\epsilon_i \in \{-1, 1\}$ $(1 \leq i \leq l)$. Then $EM = (\epsilon_i \cdot M_{ij})$ and $ME = (\epsilon_j \cdot M_{ij})$. We see that ME and EM agree on the diagonal blocks, and on the blocks where $M_{ij} = 0$. If $i > j$ and $M_{ij} \neq 0$, then there must be a vertex $T(j, q, r)$ being dominated by a vertex $T(i, q', r')$. By assumption, $\epsilon_i = \epsilon_j$. Therefore $EM = ME$ and $E \in C(Q)$.

Suppose now that $E \in \mathrm{GL}(|\Delta|, \mathbb{Z})$ fails the first condition. Without loss of generality, suppose $E(p, q, 1) \neq E(p, q', 1)$. Since, by definition, $\mathrm{Aut}(\Delta)$ acts transitively on the elements of $\mathcal{O}_{[v_p]}$, there is some $P \in \mathrm{GL}(|\Delta|, \mathbb{Z})$ induced by some $\phi \in \mathrm{Aut}(\Delta)$ which acts by exchanging the qth and q'th domination classes. A standard calculation shows that $[E, P] \neq 1$.

Finally, suppose $E \in \mathrm{GL}(|\Delta|, \mathbb{Z})$ fails the second condition. Assume that $T(p, q, r)$ is dominated by $T(p', q', r')$, but that $E(p, q, r) \neq E(p', q', r')$. In this case, E fails to centralise the elementary matrix which is the result of transvecting $T(p, q, r)$ by $T(p', q', r')$. □

3.3.5 Extending elements of $C(Q)$ to automorphisms of $\mathrm{Aut}(A_\Gamma)$

Using the map R from section 3.1.2, for $A \in C(Q) = C_1$ we obtain $R(A) \in \mathrm{Aut}(\mathrm{Aut}(A_\Gamma))$ which acts as A on $\mathbb{Z}^{k|\Delta|} \leq \mathrm{Aut}(A_\Gamma)$ and as the identity on $\mathrm{GL}(k, \mathbb{Z}) \times \mathrm{Aut}(A_\Delta) \leq \mathrm{Aut}(A_\Gamma)$.

If there are d domination classes in Δ, then $|C_1| \le 2^d$, by Proposition 3.3.6. We now determine $\hat{R}(C_1)$, the image of $R(C_1)$ in $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$.

Let $nh \in \mathbb{Z}^{k|\Delta|} \rtimes [\mathrm{GL}(k,\mathbb{Z}) \times \mathrm{Aut}(A_\Delta)]$, with $h \neq 1$. Conjugating $\mathrm{Aut}(A_\Gamma)$ by nh fixes $\mathrm{GL}(k,\mathbb{Z}) \times \mathrm{Aut}(A_\Delta)$ pointwise only if h is central in $\mathrm{GL}(k,\mathbb{Z}) \times \mathrm{Aut}(A_\Delta)$. The only such non-trivial central element is ι, the automorphism inverting each generator of $\mathbb{Z}^k$ (see Proposition 3.5.1). Given that $\alpha_\iota(n) = -n$ for each $n \in \mathbb{Z}^{k|\Delta|}$, we see that for any $m \in \mathbb{Z}^{k|\Delta|}$, we have $(m,1)^{(n,\iota)} = (-m,1)$.

So, regardless of which n we choose, the automorphism of $\mathrm{Aut}(A_\Gamma)$ induced by conjugation by $n\iota$ is equal to $R(-I_{k|\Delta|})$. In other words, when we conjugate by $n\iota$, we map each lateral transvection to its inverse. Thus, for $A, B \in C_1$, $R(AB^{-1})$ is inner if and only if $A(p,q,r) = -B(p,q,r)$ for every p, q, and r: that is, if and only if $AB^{-1} = -I_{k|\Delta|}$. This means $|R(C_1)| = 2|\hat{R}(C_1)|$.

3.3.6 First proof of Theorem 3.1.2

We are now able to prove Theorem 3.1.2 for right-angled Artin groups with non-trivial centre.

Proof (1) of Theorem 3.1.2. By Proposition 3.3.3, we have a split decomposition of $\mathrm{Aut}(A_\Gamma)$, whose kernel is $\mathbb{Z}^{k|\Delta|}$. The structure of $C_1 = C(Q)$ is given by Proposition 3.3.6. We have fewest constraints on C_1 if Δ is such that domination occurs only between vertices in the same domination class, and when each domination class lies in an $\mathrm{Aut}(\Delta)$-orbit by itself. This is achieved, for example, if $\Delta = \mathcal{C}$, a disjoint union of pairwise non-isomorphic complete graphs, each of rank at least two. The graph symmetries of $\mathcal{C}$ form a direct product of symmetric groups, as vertices may only be permuted with ones in their own connected components. Similarly, a vertex is dominated by another if and only if they belong to the same connected component. Thus, the domination equivalence classes of $\mathcal{C}$ sit in $\mathrm{Aut}(\mathcal{C})$-orbits by themselves.

Suppose $\mathcal{C}$ has d connected components. For $A \in C(Q)$, Proposition 3.3.6 implies A is entirely determined by the entries $A(p,1,1)$ $(1 \le p \le d)$, working within a fixed Δ-block. This gives $|C(Q)| = 2^d$, and so the image of $C(Q)$ in $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ has order 2^{d-1}, by the discussion in Section 3.3.5. As we may choose d to be as large as we like, the result follows. $\square$

3.4 Proof of Theorem 3.1.2: centreless right-angled Artin groups

In this section, we demonstrate that Theorem 3.1.2 also holds when working over classes of centreless right-angled Artin groups. From now on, we assume that the graph Γ has no social vertices, so that A_Γ has trivial centre. A simplicial graph $\Gamma = (V, E)$ is said to have *no separating intersection of links ('no SILs')* if for all $v, w \in V$ with v not adjacent to w, each connected component of $\Gamma \setminus (\mathrm{lk}(v) \cap \mathrm{lk}(w))$ contains either v or w. We have the following theorem.

Theorem 3.4.1 (Charney-Ruane-Stambaugh-Vijayan [15]). *Let Γ be a finite simplicial graph with no SILs. Then $\mathrm{PC}(A_\Gamma)$, the subgroup of $\mathrm{Aut}(A_\Gamma)$ generated by partial conjugations, is a right-angled Artin group, whose defining graph has vertices in bijection with the partial conjugations of A_Γ.*

We restrict ourselves to looking at certain no SILs graphs, to obtain a nice decomposition of $\mathrm{Aut}(A_\Gamma)$. We say a graph Γ is *weakly austere* if it has trivial symmetry group and no dominated vertices. Note that this is a loosening of the definition of an austere graph: removing a vertex star need no longer leave the graph connected.

Lemma 3.4.2. *Let $\Gamma = (V, E)$ be weakly austere and have no SILs. For $c \in V$, let $K_c = |\pi_0(\Gamma \setminus \mathrm{st}(c))|$. Then*

$$|\mathrm{Out}(\mathrm{Aut}(A_\Gamma))| \geq 2^{K_c - 1}.$$

Proof. Since Γ is weakly austere, the only LS generators which are defined are the inversions and the partial conjugations. Letting I_Γ denote the finite subgroup generated by the inversions ι_v $(v \in V)$, we obtain the decomposition

$$\mathrm{Aut}(A_\Gamma) \cong \mathrm{PC}(A_\Gamma) \rtimes I_\Gamma,$$

where the inversions act by inverting or commuting with partial conjugations in the obvious way. Since Γ has no SILs, it follows from Theorem 3.4.1 that $\mathrm{PC}(A_\Gamma) \cong A_\Delta$ for some simplicial graph Δ whose vertices are in bijection with the partial conjugations of A_Γ.

Fix $c \in V$ and let $\{\gamma_{c,D_i} \mid 1 \leq i \leq K_c\}$ be the set of partial conjugations by c. Let $\eta_{c,j}$ be the LS generator of $\mathrm{Aut}(A_\Delta)$ which inverts γ_{c,D_j}, but fixes the other vertex-generators of A_Δ. This extends to an automorphism of $\mathrm{Aut}(A_\Gamma)$, by specifying that I_Γ is fixed pointwise: all that needs to be checked is that the action of I_Γ on $\mathrm{PC}(A_\Gamma)$ is preserved, which is a straightforward calculation. We abuse notation, and write $\eta_{c,j} \in \mathrm{Aut}(\mathrm{Aut}(A_\Gamma))$.

If $K_c > 1$, we see $\eta_{c,j}$ is not inner. Assume $\eta_{c,j}$ is equal to conjugation by $p\kappa \in \mathrm{PC}(A_\Gamma) \rtimes I_\Gamma$. For any $\gamma \in \mathrm{PC}(A_\Gamma)$, we have $(\gamma, 1)^{(p,\kappa)} = (p\gamma^\kappa p^{-1}, 1)$. Since, by assumption, $p(\gamma_{c,D_j})^\kappa p^{-1} = \eta_{c,j}(\gamma_{c,D_j}) = {\gamma_{c,D_j}}^{-1}$, an exponent sum argument tells us that κ must act by inverting γ_{c,D_j}, and so κ must invert c in A_Γ. (We know that the exponent sum with respect to γ_{c,D_j} is well-defined, since $\mathrm{PC}(A_\Gamma)$ is a right-angled Artin group by assumption, and it is trivial to check that exponent sums are always well-defined for right-angled Artin groups). However, $\eta_{c,j}$ fixes γ_{c,D_i} for all $i \neq j$, by definition, and a similar exponent sum argument implies that κ *cannot* invert c in A_Γ. Thus, by contradiction, $\eta_{c,j}$ cannot be inner.

As above, we may choose a subset of $\{\gamma_{c,D_i}\}$ to invert, and extend this to an automorphism of $\mathrm{Aut}(A_\Gamma)$. Take two distinct such automorphisms, η_1 and η_2. Their difference $\eta_1\eta_2^{-1}$ is inner if and only if it inverts *every* element of $\{\gamma_{c,D_i}\}$. Otherwise, we would get the same contradiction as before. An elementary counting argument gives the desired lower bound of 2^{K_c-1}: we simply choose whether or not to invert each of the K_c partial conjugations $\{\gamma_{c,D_j}\}$, then note that they get identified in pairs when they are mapped into $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$. $\square$

Observe that if Γ is austere, we cannot find a vertex c with $K_c > 1$. This is the reason we loosen the definition and consider weakly austere graphs.

3.4.1 Second proof of Theorem 3.1.2

By exhibiting an infinite family of graphs over which the size of $|\{\gamma_{c,D_i}\}|$ is unbounded, applying Lemma 3.4.2 will give a second proof of Theorem 3.1.2.

Proof (2) of Theorem 3.1.2. Fix $t \in \mathbb{Z}$ with $t \geq 3$. Define $\sigma_0 = 0$ and choose $\{\sigma_1 < \ldots < \sigma_t\} \subset \mathbb{Z}^+$ subject to the conditions:

(1) For each $0 < i \leq t$, we have $\sigma_i - \sigma_{i-1} > 2$, and

(2) If $i \neq j$, then $\sigma_i - \sigma_{i-1} \neq \sigma_j - \sigma_{j-1}$.

We use the set $E := \{\sigma_i \mid 0 \leq i \leq t\}$ to construct a simplicial graph. Begin with a cycle on σ_t vertices, labelled $0, 1, \ldots, \sigma_t - 1$ in the natural way. Join one extra vertex, labelled c, to those labelled σ_i, for $0 \leq i < t$. We denote the resulting graph by Γ_E. Figure 3.3 shows an example of such a Γ_E.

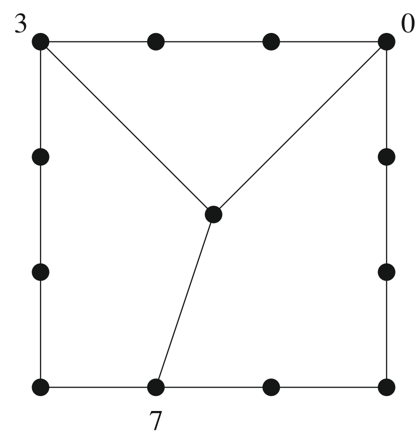

Figure 3.3: *The graph* Γ_E*, for* $E = \{3, 7, 12\}$.

For $E \subset \mathbb{Z}^+$ satisfying the above conditions, we see that Γ_E is weakly austere and has no SILs. Condition (1) ensures that no vertex is dominated by another: there is no domination in the original cycle graph we started with, so all we need to check is that no vertices dominate or are dominated by c. If c dominated some vertex $v \neq c$, then Condition (1) would be violated. Suppose v is labelled by σ_i: then its neighbours are labelled σ_{i-1}, σ_{i+1} (mod σ_t), and perhaps c. The neighbours of c are never contained in such a set, so v cannot dominate c. We note that if the inequality in Condition (1) is changed to $\sigma_i - \sigma_{i-1} \geq 2$, then domination may occur.

Observe that c is fixed by any $\phi \in \mathrm{Aut}(\Gamma_E)$, as it is the unique vertex of Γ_E whose neighbours all have three neighbours. Since each connected component of $\Gamma \setminus \mathrm{st}(c)$ has $\sigma_i - \sigma_{i-1} - 1$ elements (for some $1 \leq i \leq t$), condition (2) implies that $\mathrm{Aut}(\Gamma_E) = 1$. To see that Γ_E has no SILs, observe that the intersection of the links of any two vertices has order at most 1. When a single vertex is removed, Γ_E remains connected, and so it has no SILs.

Lemma 3.4.2 applied to the family of graphs $\{\Gamma_E\}$ proves the theorem. $\square$

3.5 Extremal behaviour and generalisations

In Sections 3.3 and 3.4, we gave examples of A_Γ for which $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ was non-trivial, but not necessarily infinite. Currently, there are very few known A_Γ for which $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ exhibits 'extremal behaviour', that is, A_Γ for which $\mathrm{Out}(\mathrm{Aut}(A_\Gamma)))$ is trivial or infinite. In this final section, we discuss the possibility of such behaviour, and generalisations of the current work to automorphism towers.

3.5.1 Complete automorphisms groups

Recall that a group G is said to be complete if it has trivial centre and every automorphism of G is inner. Our proofs of Theorems 3.1.2 and 3.1.3 relied upon us being able to exhibit large families of right-angled Artin groups whose automorphisms groups are not complete. It is worth noting that if A_Γ is not free abelian, then $\mathrm{Aut}(A_\Gamma)$ has trivial centre, so completeness of $\mathrm{Aut}(A_\Gamma)$ is not ruled out.

Proposition 3.5.1. *Let A_Γ be a right-angled Artin group. Then $Z(\mathrm{Aut}(A_\Gamma))$ has order at most two. In particular, if A_Γ is not free abelian, then $\mathrm{Aut}(A_\Gamma)$ is centreless.*

Proof. For brevity of proof, we assume that $A_\Gamma \cong \mathbb{Z}^k \times A_\Delta$, taking $k = 0$, and $\mathbb{Z}^k = 1$ if A_Γ is centreless. If A_Γ is free abelian of rank k, then $Z(\mathrm{Aut}(A_\Gamma)) \cong Z(\mathrm{GL}(k, \mathbb{Z})) \cong \mathbb{Z}/2$. From now on, we assume the centre of A_Γ is proper.

We now adapt the standard proof that a centreless group has centreless automorphism group. Suppose that $\phi \in \mathrm{Aut}(A_\Gamma)$ is central. We know that $\mathrm{Inn}(A_\Gamma) \cong A_\Gamma/\mathbb{Z}^k \cong A_\Delta$. For any $\gamma_w \in \mathrm{Inn}(A_\Gamma)$, we must have $\gamma_w = \phi\gamma_w\phi^{-1} = \gamma_{\phi(w)}$. So, for ϕ to be central, it must fix every element of A_Δ. Observe that if $k = 0$, then ϕ must be trivial, and we are done.

Assume now that $k \geq 1$. For any $\phi \in \mathrm{Aut}(A_\Gamma)$, we also have $\phi(u) \in \mathbb{Z}^k$, for all $u \in \mathbb{Z}^k$. So, a central ϕ must simply be an element of $\mathrm{GL}(k,\mathbb{Z})$, since it must be the identity on A_Δ, and take $\mathbb{Z}^k$ into itself.

In particular, we have that $Z(\mathrm{Aut}(A_\Gamma)) \leq Z(\mathrm{GL}(k,\mathbb{Z})) = \{1,\iota\}$, where ι is the automorphism inverting each generator of $\mathbb{Z}^k$. However, lateral transvections are not centralised by ι, and so the centre of $\mathrm{Aut}(A_\Gamma)$ is trivial. □

In this chapter, we have focused on finding right-angled Artin groups whose automorphism groups are not complete: an equally interesting question is which right-angled Artin groups *do* have complete automorphism groups, beyond the obvious examples of ones built out of direct products of free groups. We conjecture the following.

Conjecture 3.5.2. *When Γ is austere,* $\mathrm{Aut}(A_\Gamma)$ *is complete.*

It might also be possible to adapt Bridson-Vogtmann's geometric proof [10] of the completeness of $\mathrm{Out}(F_n)$ to find examples of A_Γ for which $\mathrm{Out}(A_\Gamma)$ is complete, using Charney-Stambaugh-Vogtmann's newly developed outer space for right-angled Artin groups [16].

3.5.2 Infinite order automorphisms

At the other extreme, we might wonder which A_Γ, if any, have $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ of infinite order. An obvious approach to this problem is to exhibit an element $\alpha \in \mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ of infinite order. The approach taken in Section 3.4, involving graphs Γ with no SILs, might seem hopeful, as we certainly know of infinite order non-inner elements of $\mathrm{Aut}(\mathrm{PC}(A_\Gamma))$: in particular, dominated transvections and partial conjugations. A key property that allowed us to extend $\eta_{c,j} \in \mathrm{Aut}(\mathrm{PC}(A_\Gamma))$ to an element of $\mathrm{Aut}(\mathrm{Aut}(A_\Gamma))$ was that it respected the natural partition of the partial conjugations by their conjugating vertex. More precisely, $\eta_{c,j}$ sent a partial conjugation by $v \in V$ to a string of partial conjugations and their inverses, each by v. This ensured that the action of I_Γ on $\mathrm{PC}(A_\Gamma)$ was preserved when we extended $\eta_{c,j}$ to be the identity on I_Γ.

It might be hoped that we could find a transvection $\tau \in \mathrm{Aut}(\mathrm{PC}(A_\Gamma))$ which also respected this partition, as τ could then easily be extended to an infinite order element of $\mathrm{Aut}(\mathrm{Aut}(A_\Gamma))$. However, it is not difficult to verify that whenever Γ has no dominated vertices, as in Section 3.4, no such τ will be well-defined. Similarly, the only obvious way to extend a partial conjugation $\gamma \in \mathrm{PC}(\mathrm{PC}(A_\Gamma))$ is to an element of $\mathrm{Inn}(\mathrm{Aut}(A_\Gamma))$. This leads us to formulate the following open question.

Question: Does there exist a simplicial graph Γ such that $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ is infinite?

It seems possible that such a Γ could exist, however the methods used in this chapter do not find one. Our main approach was to find elements of $\mathrm{Aut}(\mathrm{Aut}(A_\Gamma))$ which preserve some nice decomposition of $\mathrm{Aut}(A_\Gamma)$. To find infinite order elements of $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$, it may be necessary to loosen this constraint. This would be analogous to the situation where we find only two field automorphisms of $\mathbb{C}$ which preserve $\mathbb{R}$, but uncountably many which do not.

3.5.3 Automorphism towers

Let G be a centreless group. Then G embeds into its automorphism group, $\mathrm{Aut}(G)$, as the subgroup of inner automorphisms, $\mathrm{Inn}(G)$, and $\mathrm{Aut}(G)$ is also centreless. We inductively define

$$\mathrm{Aut}^i(G) = \mathrm{Aut}(\mathrm{Aut}^{i-1}(G))$$

for $i \geq 0$, with $\mathrm{Aut}^0(G) = G$. This yields the following chain of normal subgroups:

$$G \lhd \mathrm{Aut}(G) \lhd \mathrm{Aut}(\mathrm{Aut}(G)) \lhd \ldots \lhd \mathrm{Aut}^i(G) \lhd \ldots,$$

which we refer to as the *automorphism tower of* G. This sequence of groups is extended transfinitely using direct limits in the obvious way. An automorphism tower is said to *terminate* if there exists some i such that the canonical embedding $\mathrm{Aut}^i(G) \hookrightarrow \mathrm{Aut}^{i+1}(G)$ is an isomorphism. Observe that a complete group's automorphism tower terminates at the first step. Thomas [56] showed that any centreless group has a terminating automorphism tower, although it may not terminate after a finite number of steps. Hamkins [33] showed that the automorphism tower of *any* group terminates, although in the above definition, we have only considered automorphism towers of centreless groups. In the case of groups with non-trivial centre, we no longer have embeddings $\mathrm{Aut}^i(G) \hookrightarrow \mathrm{Aut}^{i+1}(G)$, but an analogous tower of normal subgroups may be formed using the obvious non-injective homomorphisms $\mathrm{Aut}^i(G) \to \mathrm{Aut}^{i+1}(G)$. The following problem suggests itself.

Problem: Determine the automorphism tower of A_Γ for an arbitrary Γ.

This seems a difficult problem in general: it is not even currently know what the group $\mathrm{Aut}^2(\mathrm{GL}(n,\mathbb{Z}))$ is. A first approach might be to find A_Γ for which $\mathrm{Out}(\mathrm{Aut}(A_\Gamma))$ is finite. It would then perhaps be easier to study the structure of $\mathrm{Aut}^2(A_\Gamma)$.

References

[1] Ian Agol, Daniel Groves, and Jason Manning. The virtual Haken conjecture. Pre-print, arXiv:1204.2810.

[2] M. A. Armstrong. On the fundamental group of an orbit space. *Proc. Cambridge Philos. Soc.*, 61:639–646, 1965.

[3] M. A. Armstrong. A presentation for a group of automorphisms of a simplicial complex. *Glasgow Math. J.*, 30(3):331–337, 1988.

[4] G. Baron and W. Imrich. Asymmetrische reguläre Graphen. *Acta Math. Acad. Sci. Hungar*, 20:135–142, 1969.

[5] A. Baudisch. Subgroups of semifree groups. *Acta Math. Acad. Sci. Hungar.*, 38(1-4):19–28, 1981.

[6] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math. (2)*, 135(1):1–51, 1992.

[7] Joan S. Birman. *Braids, links, and mapping class groups.* Princeton University Press, Princeton, N.J.; University of Tokyo Press, Tokyo, 1974. Annals of Mathematics Studies, No. 82.

[8] Joan S. Birman and Hugh M. Hilden. On isotopies of homeomorphisms of Riemann surfaces. *Ann. of Math. (2)*, 97:424–439, 1973.

[9] Tara Brendle, Dan Margalit, and Andrew Putman. Generators for the hyperelliptic Torelli group and the kernel of the Burau representation at t = -1. Pre-print, arXiv:1211.4018.

[10] Martin R. Bridson and Karen Vogtmann. Automorphisms of automorphism groups of free groups. *J. Algebra*, 229(2):785–792, 2000.

[11] Kenneth S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1982.

[12] Kenneth S. Brown. Presentations for groups acting on simply-connected complexes. *J. Pure Appl. Algebra*, 32(1):1–10, 1984.

[13] W. Burnside. *Theory of groups of finite order*. Dover Publications, Inc., New York, 1955. 2d ed.

[14] Ruth Charney. On the problem of homology stability for congruence subgroups. *Comm. Algebra*, 12(17-18):2081–2123, 1984.

[15] Ruth Charney, Kim Ruane, Nathaniel Stambaugh, and Anna Vijayan. The automorphism group of a graph product with no SIL. *Illinois J. Math.*, 54(1):249–262, 2010.

[16] Ruth Charney, Nathaniel Stambaugh, and Karen Vogtmann. Outer space for right-angled artin groups i. Pre-print, arXiv:1212.4791.

[17] Ruth Charney and Karen Vogtmann. Finiteness properties of automorphism groups of right-angled Artin groups. *Bull. Lond. Math. Soc.*, 41(1):94–102, 2009.

[18] Donald J. Collins. Palindromic automorphisms of free groups. In *Combinatorial and geometric group theory (Edinburgh, 1993)*, volume 204 of *London Math. Soc. Lecture Note Ser.*, pages 63–72. Cambridge Univ. Press, Cambridge, 1995.

[19] John Crisp and Luis Paris. Artin groups of type B and D. *Adv. Geom.*, 5(4):607–636, 2005.

[20] John Crisp and Bert Wiest. Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups. *Algebr. Geom. Topol.*, 4:439–472, 2004.

[21] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84(1):91–119, 1986.

[22] Michael W. Davis and Tadeusz Januszkiewicz. Right-angled Artin groups are commensurable with right-angled Coxeter groups. *J. Pure Appl. Algebra*, 153(3):229–235, 2000.

[23] Matthew Day and Andrew Putman. A Birman exact sequence for $\mathrm{Aut}(F_n)$. *Adv. Math.*, 231(1):243–275, 2012.

[24] Matthew Day and Andrew Putman. The complex of partial bases for F_n and finite generation of the Torelli subgroup of $\mathrm{Aut}(F_n)$. *Geom. Dedicata*, 164:139–153, 2013.

[25] Matthew B. Day. Peak reduction and finite presentations for automorphism groups of right-angled Artin groups. *Geom. Topol.*, 13(2):817–855, 2009.

[26] Andrew J. Duncan and Vladimir N. Remeslennikov. Automorphisms of partially commutative groups II: Combinatorial subgroups. *Internat. J. Algebra Comput.*, 22(7):1250074, 44, 2012.

[27] Joan L. Dyer and Edward Formanek. The automorphism group of a free group is complete. *J. London Math. Soc. (2)*, 11(2):181–190, 1975.

[28] Joan L. Dyer and Edna K. Grossman. The automorphism groups of the braid groups. *Amer. J. Math.*, 103(6):1151–1169, 1981.

[29] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.

[30] Robert Frucht. Graphs of degree three with a given abstract group. *Canadian J. Math.*, 1:365–378, 1949.

[31] Henry H. Glover and Craig A. Jensen. Geometry for palindromic automorphism groups of free groups. *Comment. Math. Helv.*, 75(4):644–667, 2000.

[32] Richard Hain. Finiteness and Torelli spaces. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 57–70. Amer. Math. Soc., Providence, RI, 2006.

[33] Joel David Hamkins. Every group has a terminating transfinite automorphism tower. *Proc. Amer. Math. Soc.*, 126(11):3223–3226, 1998.

[34] Susan Hermiller and John Meier. Algorithms and geometry for graph products of groups. *J. Algebra*, 171(1):230–257, 1995.

[35] L. K. Hua and I. Reiner. Automorphisms of the unimodular group. *Trans. Amer. Math. Soc.*, 71:331–348, 1951.

[36] Nikolai V. Ivanov and John D. McCarthy. On injective homomorphisms between Teichmüller modular groups. I. *Invent. Math.*, 135(2):425–486, 1999.

[37] Craig Jensen, Jon McCammond, and John Meier. The Euler characteristic of the Whitehead automorphism group of a free product. *Trans. Amer. Math. Soc.*, 359(6):2577–2595, 2007.

[38] D. G. Khramtsov. Completeness of groups of outer automorphisms of free groups. In *Group-theoretic investigations (Russian)*, pages 128–143. Akad. Nauk SSSR Ural. Otdel., Sverdlovsk, 1990.

[39] Ryoma Kobayashi. A finite presentation of the level 2 principal congruence subgroup of $GL(n;\mathbb{Z})$. Preprint in preparation, 2014.

[40] Sava Krstić and James McCool. The non-finite presentability of $\mathrm{IA}(F_3)$ and $\mathrm{GL}_2(\mathbf{Z}[t,t^{-1}])$. *Invent. Math.*, 129(3):595–606, 1997.

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

[42] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory.* Springer-Verlag, Berlin-New York, 1977. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.

[43] Wilhelm Magnus. Über n-dimensionale Gittertransformationen. *Acta Math.*, 64(1):353–367, 1935.

[44] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial group theory.* Dover Publications, Inc., New York, revised edition, 1976. Presentations of groups in terms of generators and relations.

[45] Dan Margalit. Private communication, 2013.

[46] John D. McCarthy and Ulrich Pinkall. Representing homology automorphisms of nonorientable surfaces. Max Planck Inst. preprint MPI/SFB 85-11, 2004.

[47] John Milnor. *Introduction to algebraic K-theory.* Princeton University Press, Princeton, N.J.; University of Tokyo Press, Tokyo, 1971. Annals of Mathematics Studies, No. 72.

[48] Takayuki Morifuji. On Meyer's function of hyperelliptic mapping class groups. *J. Math. Soc. Japan*, 55(1):117–129, 2003.

[49] Jakob Nielsen. Die Isomorphismengruppe der freien Gruppen. *Math. Ann.*, 91(3-4):169–209, 1924.

[50] I. B. S. Passi, Mahender Singh, and Manoj K. Yadav. Automorphisms of abelian group extensions. *J. Algebra*, 324(4):820–830, 2010.

[51] B. Perron and J. P. Vannier. Groupe de monodromie géométrique des singularités simples. *Math. Ann.*, 306(2):231–245, 1996.

[52] Andrew Putman. Obtaining presentations from group actions without making choices. *Algebr. Geom. Topol.*, 11(3):1737–1766, 2011.

[53] Jean-Pierre Serre. *Trees.* Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[54] Herman Servatius. Automorphisms of graph groups. *J. Algebra*, 126(1):34–60, 1989.

[55] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[56] Simon Thomas. The automorphism tower problem. *Proc. Amer. Math. Soc.*, 95(2):166–168, 1985.

[57] Richard D. Wade. Johnson homomorphisms and actions of higher-rank lattices on right-angled Artin groups. *J. Lond. Math. Soc. (2)*, 88(3):860–882, 2013.

[58] Richard D. Wade. Folding free-group automorphisms. *Q. J. Math.*, 65(1):291–304, 2014.

[59] Charles Wells. Automorphisms of group extensions. *Trans. Amer. Math. Soc.*, 155:189–194, 1971.

[60] Daniel T. Wise. *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics.*

Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012.