msp

Algebraic & Geometric Topology 15 (2015) 3535–3567

A generating set for the palindromic Torelli group

NEIL J FULLARTON

A *palindrome* in a free group F_n is a word on some fixed free basis of F_n that reads the same backwards as forwards. The *palindromic automorphism group* ΠA_n of the free group F_n consists of automorphisms that take each member of some fixed free basis of F_n to a palindrome; the group ΠA_n has close connections with hyperelliptic mapping class groups, braid groups, congruence subgroups of $GL(n, \mathbb{Z})$, and symmetric automorphisms of free groups. We obtain a generating set for the subgroup of ΠA_n consisting of those elements that act trivially on the abelianisation of F_n, the *palindromic Torelli group* $\mathcal{PI}_n$. The group $\mathcal{PI}_n$ is a free group analogue of the hyperelliptic Torelli subgroup of the mapping class group of an oriented surface. We obtain our generating set by constructing a simplicial complex on which $\mathcal{PI}_n$ acts in a nice manner, adapting a proof of Day and Putman. The generating set leads to a finite presentation of the principal level 2 congruence subgroup of $GL(n, \mathbb{Z})$.

20F65, 57M07, 57MXX

1 Introduction

Let F_n be the free group of rank n on some fixed free basis X. The palindromic automorphism group of F_n, denoted ΠA_n, consists of automorphisms of F_n that take each member of X to some palindrome, that is, a word on X that reads the same backwards as forwards. Collins [8] introduced the group ΠA_n and proved that it is finitely presented, giving an explicit presentation. Glover and Jensen [15] obtained further results about ΠA_n, utilising a contractible subspace of the auter space of F_n on which ΠA_n acts cocompactly, with finite stabilisers. For instance, they calculate that the virtual cohomological dimension of ΠA_n is $n-1$. The group ΠA_n is a free group analogue of the hyperelliptic mapping class group of an oriented surface; we develop this analogy later in this introduction.

In this paper, we are primarily concerned with the intersection of ΠA_n with the Torelli subgroup of F_n, that is, the subgroup of automorphisms of ΠA_n that act trivially on the abelianisation of F_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 [8] first observed that it is non-trivial, and Jensen, McCammond and Meier

Published: 12 January 2016 DOI: 10.2140/agt.2015.15.3535

[17, Corollary 6.3] showed that $\mathcal{PI}_n$ is not of finite homological type for $n \geq 3$. In Section 2, we introduce non-trivial members of $\mathcal{PI}_n$ $(n \geq 3)$ known as *doubled commutator transvections* and *separating* π*–twists*. The main theorem of this paper establishes that these generate $\mathcal{PI}_n$.

Theorem A *The group* $\mathcal{PI}_n$ $(n \geq 3)$ *is generated by doubled commutator transvections and separating* π*–twists.*

In order to prove Theorem A, we establish finite generating sets for the subgroups of $\Pi\mathrm{A}_n$ consisting of automorphisms that fix each member of some specified subset of the free basis X. These generating sets, which are given precisely in the statement of Proposition 2.2, are obtained by utilising Stallings' graph folding algorithm.

Let $\Gamma_n[2]$ denote the *principal level* 2 *congruence subgroup of* $\mathrm{GL}(n,\mathbb{Z})$, that is, the kernel of the surjection $\mathrm{GL}(n,\mathbb{Z}) \to \mathrm{GL}(n,\mathbb{Z}/2)$ that reduces matrix entries mod 2. In Section 2, we discuss a short exact sequence with kernel the palindromic Torelli group and quotient $\Gamma_n[2]$. For $1 \leq i \neq j \leq n$, let $S_{ij} \in \Gamma_n[2]$ be the matrix that has 1s on the diagonal and 2 in the (i,j) position, with 0s elsewhere, and let $O_i \in \Gamma_n[2]$ differ from the identity only in having -1 in the (i,i) position. The following corollary of Theorem A provides a finite presentation of $\Gamma_n[2]$ for $n \geq 4$.

Corollary 1.1 *The principal level* 2 *congruence group* $\Gamma_n[2]$ $(n \geq 4)$ *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 different.*

We note that in the proof of Theorem A 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).

We also derive the following similar presentation for $\Gamma_n[2]$ when $n = 2$ or 3; however, these are acquired independently of Theorem A. Indeed, the presentation of $\Gamma_3[2]$ is used to obtain a generating set for $\mathcal{PI}_3$, which forms the base case of an inductive proof of Theorem A.

Proposition 1.2 *The principal level* 2 *congruence group* $\Gamma_n[2]$ $(n=2,3)$ *is generated by*

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

subject to the defining relators in the statement of Corollary 1.1 of types

- (1)–(4) *for* $n = 2$,
- (1)–(6), (8)–(10) *for* $n = 3$.

A key tool in the proof of Proposition 1.2 is an "even" version of the division algorithm for the integers. This is the observation that under certain circumstances, the quotient $q \in \mathbb{Z}$ given when dividing $a \in \mathbb{Z}$ by $b \in \mathbb{Z}$ may be chosen to be even, if we sacrifice control of the sign of the remainder $r \in \mathbb{Z}$. More details of this procedure are given in the proofs of Lemma 2.4 and Theorem 5.1.

A similar presentation for $\Gamma_n[2]$ was recently found independently by Kobayashi [18], and was also known to Margalit and Putman. As Margalit and Putman pointed out, this is a natural presentation for $\Gamma_n[2]$, as relators of types (6)–(9) bear a strong resemblance to the Steinberg relations that hold between the transvections generating $\mathrm{SL}(n, \mathbb{Z})$; see Milnor [22, Section 5].

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 subgroup of the mapping class group of an oriented surface. Let S_g and S_g^1 denote the compact, connected, oriented surfaces of genus g with 0 and 1 boundary components, respectively. We shall use S to denote such a surface, with or without boundary. Recall that the *mapping class group* of the surface S, denoted $\mathrm{Mod}(S)$, consists of isotopy classes of orientation-preserving self-homeomorphisms of S, where isotopies are required to fix any boundary component pointwise at all times. For a self-homeomorphism f of S, we denote its isotopy class by $[f]$.

A *hyperelliptic involution* of the surface S is an order-2 homeomorphism of the surface that acts as $-I$ on $H_1(S, \mathbb{Z})$; see Brendle and Margalit [4, Sections 2 & 4]. Let s denote the homeomorphism of S_g^1 seen in Figure 1. By capping the boundary with a disk, the map s induces a homeomorphism of S_g, which we also denote s, by an abuse of notation. The map s is an example of a hyperelliptic involution of S_g^1 (and S_g). We note that the mapping class of any hyperelliptic involution in $\mathrm{Mod}(S_g)$ $(g \geq 1)$ is conjugate to $[s]$; see Farb and Margalit [12, Proposition 7.15].

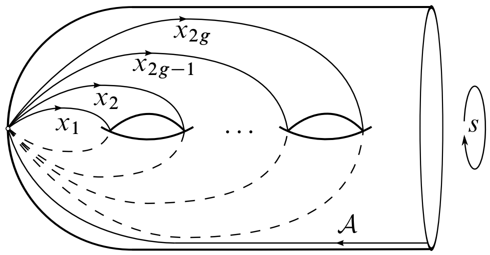

Figure 1: The involution s rotates the surface by π radians. Under the Nielsen embedding, we may view the braid group $B_{2g} \leq \mathrm{SMod}(S_g^1)$ as a subgroup of $\Pi\mathrm{A}_{2g} \leq \mathrm{Aut}(F_{2g})$.

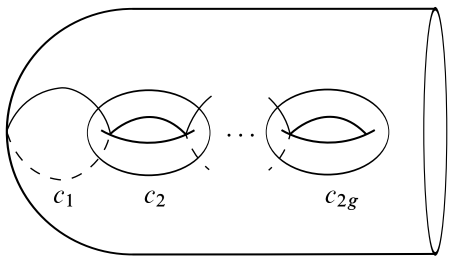

Figure 2: The standard symmetric chain in S_g^1. The Dehn twists about $c_1, \dots, c_{2g}$ generate $\mathrm{SMod}(S_g^1) \cong B_{2g+1}$.

The *hyperelliptic mapping class group of the surface* S_g, denoted $\mathrm{SMod}(S_g)$, is the centraliser of $[s]$ in $\mathrm{Mod}(S_g)$. Although $[s] \notin \mathrm{Mod}(S_g^1)$, as s does not fix the boundary of S_g^1, we define the *hyperelliptic mapping class group of* S_g^1, denoted $\mathrm{SMod}(S_g^1)$, to be the group of isotopy classes of the centraliser of s in $\mathrm{Homeo}^+(S_g^1)$ [12, Chapter 9].

An obvious analogue of a hyperelliptic involution in $\mathrm{Aut}(F_n)$ is an order-2 member of $\mathrm{Aut}(F_n)$ that acts as $-I$ on $H_1(F_n, \mathbb{Z}) = \mathbb{Z}^n$. An example of such an involution in $\mathrm{Aut}(F_n)$ is the automorphism ι that inverts each member of the free basis X. An analogy between s and ι is strengthened by two observations. Firstly, Glover and Jensen [15, Proposition 2.4] showed that any hyperelliptic involution in $\mathrm{Aut}(F_n)$ is conjugate to ι. Secondly, the action of s on $\pi_1(S_g^1) = F_{2g}$, with free basis as seen in Figure 1, is to invert each member of the free basis, as ι does. It is easily verified that $\Pi\mathrm{A}_n$ is the centraliser of ι in $\mathrm{Aut}(F_n)$ [15, Section 2], so we may think of $\Pi\mathrm{A}_n$ as being a free group analogue of the hyperelliptic mapping class groups $\mathrm{SMod}(S_g)$ and $\mathrm{SMod}(S_g^1)$.

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 1 as a free basis for $\pi_1(S_g^1)$. Observe that s acts on these loops by switching

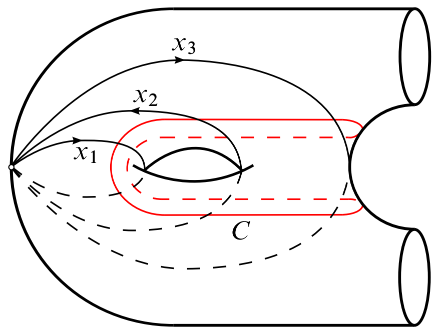

Figure 3: The Dehn twist about the symmetric, separating curve C maps to a separating π–twist in $\mathcal{PI}_{2g}$ under the Nielsen embedding. Note that we only depict a genus-one subsurface of S_g^1, and that x_2 has a different orientation than in Figure 1.

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 [3], and is generated by Dehn twists about the curves in the standard, symmetric chain on S_g^1, seen in Figure 2. 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 1 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} is taken to the automorphism Q_i of the form

$$x_i \mapsto x_{i+1}, \quad x_{i+1} \mapsto x_{i+1} x_i{}^{-1} x_{i+1}, \quad x_j \mapsto x_j$$

for $1 \le i < 2g$ and $j \ne i, i+1$. 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 and Vannier [24] and Crisp and Paris [9]. 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 $\Pi\mathrm{A}_{2g-1} \le \mathrm{Aut}(F_{2g})$.

Palindromic and hyperelliptic Torelli groups The main focus of our study in this paper 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 $\mathcal{SI}_g^1 := \mathrm{SMod}(S_g^1) \cap \mathcal{I}_g^1$ to be the *hyperelliptic Torelli group*. Brendle, Margalit and Putman [5] recently proved a conjecture of Hain [16], also stated by Morifuji [23], showing that $\mathcal{SI}_g^1$ is generated by Dehn twists about separating simple closed curves of genus one and two that are fixed by s.

Our generating set for $\mathcal{PI}_n$ compares favourably with Brendle, Margalit and Putman's for $\mathcal{SI}_g^1$, in the following way. We shall see in Section 2 that any Dehn twist about a symmetric separating curve of genus one that lies in the pre-image of the Nielsen embedding discussed above, maps to a separating π–twist in $\mathcal{PI}_{2g}$. In fact, up to conjugation by ΠA_{2g}, this is the definition of a separating π–twist. The Dehn twist about the curve C shown in Figure 3 is an example of such a mapping class. Note that the Dehn twist about C is one of the generators of Brendle, Margalit and Putman. We shall see in Proposition 3.7 that doubled commutator transvections do not suffice to generate $\mathcal{PI}_n$, so we observe that our generating set involves Brendle, Margalit and Putman's generators in a significant way. Thus, 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.

One way in which the analogy between $\mathcal{PI}_n$ and $\mathcal{SI}_g^1$ breaks down, however, is their behaviour when ΠA_n and $\mathrm{SMod}(S_g^1)$ are abelianised, to $(\mathbb{Z}/2)^3$ and $\mathbb{Z}$, respectively. An immediate corollary of Theorem A is that $\mathcal{PI}_n$ vanishes in the abelianisation of ΠA_n. In contrast, the image of $\mathcal{SI}_g^1$ in the abelianisation of $\mathrm{SMod}(S_g^1)$ is $4\mathbb{Z}$, which may be shown by calculating the images of Brendle, Margalit and Putman's generators in the abelianisation of $\mathrm{SMod}(S_g^1)$.

Palindromes in right-angled Artin groups In forthcoming work with Anne Thomas [14], we extend Collins' definition of palindromic automorphisms to the right-angled Artin group setting. We obtain generating sets for the analogously defined palindromic automorphism group and palindromic Torelli group of an arbitrary right-angled Artin group.

Approach of the paper To prove Theorem A, we employ a standard, geometric technique: we find a sufficiently connected complex on which $\mathcal{PI}_n$ acts with sufficiently connected quotient, and use a theorem of Armstrong [1] to conclude that $\mathcal{PI}_n$ is generated by the action's vertex stabilisers. This approach is modelled on a proof of Day and Putman [11], which recovers Magnus' finite generating set for the Torelli subgroup of $\mathrm{Aut}(F_n)$.

Conventions We apply functions from right to left. For $g, h \in G$ a group, we let $[g,h] = ghg^{-1}h^{-1}$. In a graph, we denote an edge between vertices x and y by $x - y$. In a group G, we will also conflate a relation $P = Q$ with the relator PQ^{-1} when this is unambiguous.

Outline of the paper In Section 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 3, we introduce the complex of partial π–bases of F_n, and use it to obtain a generating set for $\mathcal{PI}_n$. In Section 4, we prove key results about the connectivity of the complexes involved in the proof of Theorem A. In Section 5, we obtain a finite presentation for $\Gamma_3[2]$ used in the base case of our inductive proof of Theorem A.

Acknowledgements The results in this paper formed part of the author's PhD thesis [13]. The author thanks his PhD supervisor Tara Brendle for her guidance, and for introducing him to the palindromic Torelli group. The author is grateful for the hospitality of the Institute for Mathematical Research at Eidgenössische Technische Hochschule Zürich, where part of this work was completed. The author also thanks Ruth Charney and Karen Vogtmann for helpful discussions, and Thomas Church, Dan Margalit, Luis Paris, Andrew Putman, Richard Wade and Liam Watson for useful comments on an earlier version of this paper. The author is grateful to the referee for numerous constructive comments.

2 The palindromic automorphism group

Let F_n be the free group of rank n, on some fixed free basis $X := \{x_1, \ldots, x_n\}$. For a word $w = l_1 \cdots l_k$ on $X^{\pm 1}$, let w^{rev} denote the *reverse* of w; that is, we have $w^{\mathrm{rev}} = l_k \cdots l_1$. 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 automorphism $\alpha \in \mathrm{Aut}(F_n)$ is said to be *palindromic* (with respect to the fixed free basis X) 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. The following proposition gives us information about the form of the palindromes $\alpha(x_i)$.

Proposition 2.1 *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 $p = w^{\mathrm{rev}} x_i^{\epsilon_i} w \in F_n$ of odd length ($w \in F_n$, $x_i \in X$, $\epsilon_i \in \{\pm 1\}$), let $c(p) = x_i$. The following argument is implicit in the work of Collins [8].

Let $\alpha \in \Pi\mathrm{A}_n$. Since $\alpha(X)$ is a free basis, its image under the natural surjection $F_n \to (\mathbb{Z}/2)^n$ must suffice to generate $(\mathbb{Z}/2)^n$. If some $\alpha(x_i)$ is of even length, it will have zero image, and so the image of $\alpha(X)$ could not generate $(\mathbb{Z}/2)^n$. If

$c(\alpha(x_i)) = c(\alpha(x_j))$ for some $i \neq j$, then $\alpha(x_i)$ and $\alpha(x_j)$ will have the same image in $(\mathbb{Z}/2)^n$, and so again $\alpha(X)$ could not generate $\alpha(\mathbb{Z}/2)^n$. □

Finite generation of ΠA_n Collins first studied the group ΠA_n, giving a finite presentation for it. For $i \neq j$, let $P_{ij} \in \Pi A_n$ 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 $\iota_j \in \Pi A_n$ 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 $P\Pi A_n$.

Collins showed that $\Pi A_n \cong E\Pi A_n \rtimes \Omega^{\pm 1}(X)$ for $n \geq 2$, where $E\Pi A_n = \langle P_{ij} \rangle$. The group $\Omega^{\pm 1}(X)$ acts on $E\Pi A_n$ in the natural way, and a defining set of relations for $E\Pi A_n$ is given by

(1) $[P_{ik}, P_{jk}] = 1$,

(2) $[P_{ij}, P_{kl}] = 1$,

(3) $P_{ij} P_{jk} P_{ik} = P_{ik}^{-1} P_{jk} P_{ij}$,

where i, j, k, l are pairwise different and the obviously undefined relators are omitted in the $n = 2$ and $n = 3$ cases.

We remark that, as noted by Collins [8], this presentation of $E\Pi A_n$ is very similar to one given for the *pure symmetric automorphism group of* F_n, $P\Sigma A_n$, which consists of automorphisms taking each $x \in X$ to a conjugate of itself. This similarity is not entirely surprising, as we may think of a palindrome yxy as a conjugate yxy^{-1}, working "mod 2" ($x, y \in X$). The embedding $B_n \hookrightarrow \Pi A_n$ discussed in Section 1 bears a striking resemblance to Artin's faithful representation of B_n into ΣA_n, the full symmetric automorphism group, whose members take each $x \in X$ to *some* conjugate [2, Corollary 1.8.3]; this similarity arises via the branched double cover map $S_g^1 \to D_{2g+1}$ [12, Figure 9.13].

Using graph folding techniques of Stallings, we obtain a new proof of finite generation of ΠA_n, as well as finding generating sets for certain fixed-point subgroups of ΠA_n. We first introduce the notation and terminology of Wade [26] regarding graph folding.

Let R_n denote the wedge of n copies of S^1 at a point o. We canonically identify $\pi_1(R_n, o)$ with F_n by selecting an orientation of each S^1, and labelling the i^{th} copy of S^1 by $x_i \in X$. We shall let $\overline{x}_i$ denote the edge obtained by reversing the orientation of x_i.

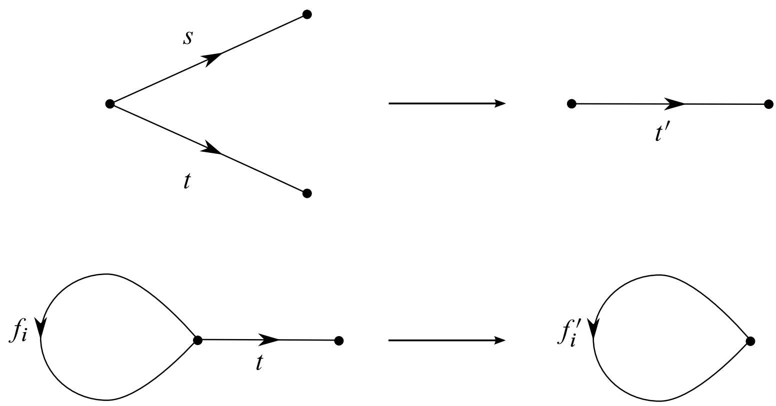

Figure 4: The two types of folding that may occur for our graph morphism ϕ. Wade [26] refers to the top fold as a type 1 fold, and to the bottom as a type 2 fold. The edges are labelled suggestively: we will demand that $s, t \in T$ and $f_i \notin T$.

Now, let Y be a finite graph of rank n with basepoint b. We will view our graphs as combinatorial objects, rather than topological ones. In particular, morphisms between graphs must take edges to edges, rather than edge-paths. A free basis for the (free) fundamental group $\pi_1(Y, b)$ is obtained in the usual way, by selecting a maximal tree T in Y, then choosing an orientation of the edges $f_1, \dots, f_n$ in Y but not T. To be consistent with Wade, we canonically orient an edge e of T by declaring its initial vertex $i(e)$ to be the one closer to the basepoint b under the edge-path metric on T.

Suppose $\theta\colon Y \to R_n$ is a morphism of graphs that induces an isomorphism of fundamental groups. The morphism θ, together with the choice of basepoint b, maximal tree T and an ordering L of the (oriented) edges of $Y \setminus T$ form a *branding* of the graph Y. A graph Y together with a 4-tuple $\mathcal{G} = (b, T, L, \theta)$ form a *branded graph* with *branding* $\mathcal{G}$.

Each branded graph Y with branding $\mathcal{G} = (b, T, L, \theta)$ yields an automorphism $B_{\mathcal{G}} \in \mathrm{Aut}(F_n)$, as follows. For each x_i in the free basis X of F_n, we have

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

where $\{y_1, \dots, y_n\}$ is the free basis of $\pi_1(Y, b)$ arising from the choices of b, T and L in the branding $\mathcal{G}$, and $\theta_*\colon \pi_1(Y, b) \to \pi_1(R_n, o)$ is the map induced by θ.

If the morphism θ maps a pair of edges e_1 and e_2 with $i(e_1) = i(e_2)$ to the same edge l of R_n, then θ factors through the quotient graph Y' of Y obtained by *folding* e_1 and e_2 together: that is, the graph obtained by identifying e_1 with e_2, and also their terminal vertices, $t(e_1)$ and $t(e_2)$, with each other. In particular, if $q\colon Y \to Y'$ is the quotient map obtained by the folding, then there is a unique graph morphism $\theta'\colon Y' \to R_n$ such that $\theta = \theta' \circ q$. While Stallings considered more general foldings, since we require θ to induce an isomorphism of fundamental groups, only two types of folding may arise for us, which are shown in Figure 4.

If we insist that the edges s and t seen in Figure 4 lie in T, and that the edge f_i does not, carrying out either type of 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). 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 the maximal tree so that the relevant edges lie in the new tree. This defines a new branding $\mathcal{G}''$ of Y. In either case, it may be shown via a careful consideration of $\pi_1(Y, b)$ (see [26, Propositions 3.2 and 3.3]) that $B_{\mathcal{G}} = B_{\mathcal{G}'} \cdot W'$ and $B_{\mathcal{G}} = B_{\mathcal{G}''} \cdot W''$, where W' and W'' are specified *Whitehead automorphisms* of F_n. These are automorphisms which fix some $x \in X$ and send each $x_i \in X \setminus \{x\}$ to one of x_i, $x_i x^{\epsilon_i}$, $x^{\epsilon_i} x_i$ or $x^{\epsilon_i} x_i x^{-\epsilon_i}$ for some $\epsilon_i \in \{\pm 1\}$.

Stallings' folding algorithm allows us to repeatedly fold the graph Y and its quotients, beginning with the morphism $\theta\colon Y \to R_n$, then continuing to fold via $\theta'\colon Y' \to R_n$, and so on. This procedure eventually terminates when we exhaust the edges we are able to fold; in this case, Stallings showed that the quotient graph is R_n, and so the morphism $\psi\colon R_n \to R_n$ obtained by repeatedly folding via θ simply permutes and perhaps inverts the n loops in R_n. This folding procedure allows us to write the automorphism $B_{\mathcal{G}}$ we began with as a product of Whitehead automorphisms and permutations and inversions of X.

With the details of folding established, we now put the algorithm to use to find generators for $\Pi\mathrm{A}_n$.

Proposition 2.2 *Fix $0 \leq k \leq n$, and let $\Pi\mathrm{A}_n(k)$ consist of automorphisms which fix $x_1, \dots, x_k$. (Our convention is 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 [26, Theorem 4.1].

We begin by introducing some terminology. Let $\phi\colon S \to T$ be an isomorphism of finite trees. For a vertex or 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 5) is the graph formed by gluing S to T along v and v', then, for each vertex $r \in S$, adding some number of edges (possibly zero) between r and r' (we allow $r = v$). We refer

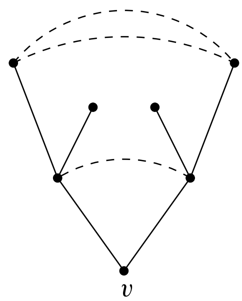

Figure 5: An example of an arch, with base point v. The dashed edges indicate the bridges that have been added to the trees that were glued together at the base point.

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 a *wedge of arches* we mean a collection of arches glued together at their base points. Note that each of the trees S_i and T_i of each arch sit inside Y as subgraphs, and Y is the union of these subgraphs, together with any bridges inside each arch.

Let $\theta\colon Y \to R_n$ be a graph morphism, with Y a wedge of arches. We call θ *symmetric* if for each edge s_i in each tree S_i in each arch of Y we have $\theta(s_i') = \theta(\overline{s_i})$. We shall define two new types of folding that we may carry out to any symmetric morphism $\theta\colon Y \to R_n$, with the resulting morphism $\theta'\colon Y' \to R_n$ on the folded graph Y' also being symmetric.

Let $\alpha \in \Pi\mathrm{A}_n(k)$. We may realise α as a morphism of graphs $\theta\colon Z \to R_n$, where Z 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 i^{th} copy of S^1. Precisely, the j^{th} edge of the oriented, subdivided S^1 corresponding to $\alpha(x_i)$ is mapped to the loop in R_n corresponding to the j^{th} letter of $\alpha(x_i)$, correctly oriented. Note that Z is a wedge of arches, and θ is symmetric by construction. We thus have $\alpha = B_{\mathcal{G}}$, where $\mathcal{G}$ is the branding of Z arising from the maximal tree that excludes the (appropriately ordered) middle subdivided edge of each copy of S^1. We now use graph folding to write α as a product of permutations, inversions and elementary palindromic automorphisms.

Let $\theta\colon Y \to R_n$ be symmetric, for some wedge of arches Y, built out of trees S_i, T_i $(1 \le i \le k)$. Since θ is symmetric, foldings of Y 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 $\theta(h_i) = \theta(h_j)$ (allowing $i = j$), then we will also be able to fold together h_i' and h_j', as they will also both have the same image under θ, namely $\theta(\overline{h_i}) = \theta(\overline{h_j})$. We call this pair of folds a *type A 2–fold*.

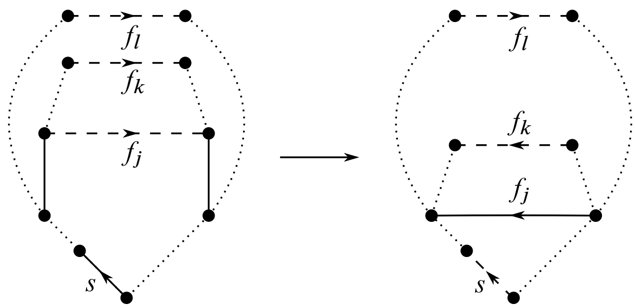

Figure 6: The two adjacent solid edges are folded onto f_j. The dashed edges represent edges excluded from the graph's chosen maximal tree. In order to record what effect this type B 2–fold has on the branded graph's associated automorphism, we must swap f_j into the maximal tree, in place of the stem s.

We may also have a sequence of edges (h_{j-1}, h_j, h_{j+1}) mapped under θ to the sequence $(\overline{x}, x, \overline{x})$ where x is an oriented edge of R_n, $h_{j-1} \in S_i$, $h_{j+1} = h'_{j-1}$ and h_j is a bridge. 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 6.

Doing either of these 2–folds to Y yields another, different wedge of arches, Y', say. A type B 2–fold simply removes an edge of valence one from S_i (and its corresponding edge in T_i) by folding it onto a bridge, producing new trees S'_i and T'_i which we use to construct Y' as a wedge of arches. A type A 2–fold similarly alters the trees S_i, S_j, T_i and T_j, producing new trees S'_i and T'_i in a description of Y' as a wedge of arches. The morphism $\theta'\colon Y' \to R_n$ induced by the folding of Y is again symmetric: any edges s_i and s'_i that were not folded still satisfy $\theta'(s'_i) = \theta'(\overline{s_i})$ by construction of θ', but so do the images of any folded edges, given how we decompose Y' as a wedge of arches using the new trees S'_i and T'_i.

In order to see what effect these 2–folds have on $\alpha \in \Pi\mathrm{A}_n$, we must keep track of a preferred maximal tree T we define on each wedge of arches 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 (seen in Figure 6) into the maximal tree. Let $p_{i(f_j)}$ denote 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 $p_{i(f_j)}$. Using calculations of Wade [26, Propositions 3.2 and 3.3], it is straightforward to verify that the effect of swapping maximal trees in this way, 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. Precisely, let $\theta\colon Y_1 \to R_n$ be a symmetric morphism of graphs, where Y_1 has branding $\mathcal{G}_1$ and let $\mathcal{G}_2$ be the induced branding of the graph Y_2 obtained by carrying out the above series of tree swaps and folds. Then

$$\phi_{\mathcal{G}_1} = \phi_{\mathcal{G}_2} \cdot P,$$

where $\phi_{\mathcal{G}_i}$ is the automorphism of F_n associated to $\mathcal{G}_i$ $(i = 1, 2)$ and P is a product of elementary palindromic automorphisms.

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 do one of two things. If the bridge is a loop at the base point v, we carry out two type 2 folds. Otherwise, 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. As before, the automorphism of F_n associated to these sequences of steps is 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 folding algorithm by applying the appropriate automorphism from $\Omega^{\pm 1}(X)$. Since $\alpha \in \Pi\mathrm{A}_n(k)$, the graph Z 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 Z. Thus, while folding such a graph Y, we only need Collins' generators that fix the first k members of the free basis X. The proposition is thus proved. □

Corollary 2.3 *The group* $\mathrm{P}\Pi\mathrm{A}_n(k)$ *of pure palindromic automorphisms that fix* $x_1, \dots, x_k$ $(0 \leq k \leq n)$ *is generated by the set* $\{P_{ij}, \iota_i \mid i > k\}$.

The principal level 2 congruence subgroup of GL($n, \mathbb{Z}$) Recall that $\Gamma_n[2]$ denotes the *principal level* 2 *congruence subgroup* of $\mathrm{GL}(n, \mathbb{Z})$, that is, the kernel of the map $\mathrm{GL}(n, \mathbb{Z}) \to \mathrm{GL}(n, \mathbb{Z}/2)$ given by reducing matrix entries mod 2. 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 a -1 in the (i, i) position. The following lemma verifies a well-known generating set for $\Gamma_n[2]$ (see, for example, McCarthy and Pinkall [21, Corollary 2.3]). We include a proof here to introduce the idea of an "even division algorithm", which we utilise in the proof of Theorem 5.1.

Lemma 2.4 *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 i^{th} row of u. Let v_1 be the standard column vector with a 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 claim is obvious. Now suppose $|u^{(1)}| > 1$. As in the proof of Proposition 2.1, 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 $0 \le r < |u^{(1)}|$. 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 (new) j^{th} 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 $M \in \Gamma_n[2]$ to 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 $\{S_{ij}, O_i \mid i, j > 1\}$. Then we simply use even row operations to fix the top row, and finish by applying O_1 if necessary. □

By Lemma 2.4, the restriction of the canonical map $\mathrm{Aut}(F_n) \to \mathrm{GL}(n, \mathbb{Z})$ gives the short exact sequence

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

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

The rest of the paper is concerned with finding a generating set for the palindromic Torelli group $\mathcal{PI}_n$. In order to describe our generating set, we introduce some terminology.

Let Y be the image of the free basis X under some automorphism $\alpha \in \Pi\mathrm{A}_n$. The set Y is also a free basis for F_n, whose members are palindromes on X; thus, we refer to Y as a *π–basis*. An automorphism $\phi \in \mathcal{PI}_n$ is a *doubled commutator transvection* if, for some y_1, y_2, y_3 in some π–basis Y, ϕ maps y_1 to $[y_2, y_3]^{\text{rev}} y_1 [y_2, y_3]$, and fixes the other members of Y. Observe that $\phi \in \mathcal{PI}_n$ is a doubled commutator transvection if and only if ϕ is conjugate in $\Pi\mathrm{A}_n$ to the commutator $\chi_1 := [P_{12}, P_{13}]$.

An automorphism $\phi \in \mathcal{PI}_n$ is a *separating π–twist* if, for some y_1, y_2, y_3 in some π–basis Y, ϕ is given by

$$\phi(y_i) = \begin{cases} d^{\mathrm{rev}} y_1 d & \text{if } i = 1, \\ d^{-1} y_2 (d^{\mathrm{rev}})^{-1} & \text{if } i = 2, \\ d^{\mathrm{rev}} y_3 d & \text{if } i = 3, \\ y_i & \text{otherwise,} \end{cases}$$

where $d = y_1{}^{-1} y_2{}^{-1} y_3{}^{-1} y_1 y_2 y_3 \in F_n$. It is a straightforward, if lengthy, calculation to verify that $\phi \in \mathcal{PI}_n$ is a separating π–twist if and only if ϕ is conjugate in $\Pi\mathrm{A}_n$ to the automorphism

$$\chi_2 := (P_{23} P_{13}{}^{-1} P_{31} P_{32} P_{12} P_{21}{}^{-1})^2 \in \mathcal{PI}_n.$$

The definition of a separating π–twist may seem unwieldy; however, it belies a hidden geometry. The automorphism χ_2 is the image in $\mathcal{PI}_n$ under the Nielsen embedding of the Dehn twist about the curve C seen in Figure 3. We call such automorphisms separating π–twists to reflect this geometric interpretation.

Theorem A states that doubled commutator transvections and separating π–twists suffice to generate $\mathcal{PI}_n$. To prove this, we construct a new complex on which $\mathcal{PI}_n$ acts in a suitable way. We then apply a theorem of Armstrong [1] 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 A.

3 The complex of partial π–bases

Day and Putman [11] 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 after $\mathcal{B}_n$, and follow their approach to find a generating set for $\mathcal{PI}_n$.

Fix $X := \{x_1, \dots, x_n\}$ as a free basis of F_n. A *π–basis*, as discussed above, 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 $\{w_1, \dots, w_k\}$. We postpone until Section 4 the proof of the following theorem on the connectedness of $\mathfrak{B}_n^\pi$.

Theorem 3.1 *For $n \geq 3$, the complex $\mathfrak{B}_n^\pi$ is simply connected.*

Our complex $\mathfrak{B}_n^\pi$ is 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 remove this technicality, as it can be shown that two odd-length palindromes are conjugate if and only if they are equal. Given this, 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 ΠA_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 s of $\mathfrak{B}_n^\pi$, then it fixes s pointwise. Following work of Charney [7] on related complexes, we obtain that the quotient of $\mathfrak{B}_n^\pi$ by $\mathcal{PI}_n$ is highly connected.

Theorem 3.2 *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 4.

Theorems 3.1 and 3.2 allow us to apply the following theorem of Armstrong [1] to the action of $\mathcal{PI}_n$ on $\mathfrak{B}_n^\pi$, for $n \geq 4$. The statement of the theorem is as given by Day and Putman [11].

Theorem 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 equality follows from the fact that $\mathrm{IA}_2 = \mathrm{Inn}(F_2)$ and $\mathrm{Inn}(F_n) \cap \Pi A_n = 1$ 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 5.

A Birman exact sequence We require a version of the free group analogue of the Birman exact sequence, as developed by Day and Putman [10]. Recall that $\mathrm{P}\Pi A_n(k)$ consists of the pure palindromic automorphisms fixing $x_1, \ldots, x_k$.

Proposition 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 $\theta_*\colon \mathrm{P}\Pi A_n(k) \to \mathrm{P}\Pi A_{n-k}$ is induced by the map $\theta\colon F_n \to F_{n-k}$ that trivialises each $x_1, \ldots, x_k$. Let $\{y_{k+1}, \ldots, y_n\}$ be a free basis for F_{n-k}, where $\theta(x_i) = y_i$ for $k+1 \leq i \leq n$. Denote by Q_{ij} and η_i the elementary palindromic automorphism sending y_i to $y_j y_i y_j$ and the inversion sending y_i to $y_i{}^{-1}$, respectively $(k+1 \leq i \neq j \leq n)$.

By Corollary 2.3, we know that $\mathrm{P}\Pi\mathrm{A}_n(k)$ is generated by the set

$$\mathcal{S} := \{P_{ij}, \iota_i \mid i > k, 1 \le j \le n\}.$$

If $j \le k$, then $\theta_*(P_{ij})$ is trivial. If $i, j \ge k+1$, then $\theta_*(P_{ij}) = Q_{ij}$ and $\theta_*(\iota_i) = \eta_i$, so we have that θ_* is surjective, by examining Collins' generators for $\mathrm{P}\Pi\mathrm{A}_{n-k}$. Indeed, the map θ_* has a section, taking Q_{ij} to P_{ij} and η_i to ι_i, which we know is well-defined by Collins' finite presentation for $\mathrm{P}\Pi\mathrm{A}_{n-k}$. Thus, we obtain a split short exact sequence via the epimorphism θ_*.

All that is left to establish is the kernel of θ_*. Notice that we have a presentation for $\mathrm{P}\Pi\mathrm{A}_{n-k}$ in terms of the generating set $\theta_*(\mathcal{S})$: explicitly, we add the relations $\theta_*(P_{ij}) = 1$ for $j \le k$ to Collins' relations on the set $\{Q_{ij}, \eta_i\}$. It is a standard fact (see, for example, Magnus, Karrass and Solitar [20, proof of Theorem 2.1]) that the kernel of θ_* is the normal closure in $\mathrm{P}\Pi\mathrm{A}_n(k)$ of the obvious lifts of the defining relators on $\theta_*(\mathcal{S})$. The only defining relators with non-trivial lifts in $\mathrm{P}\Pi\mathrm{A}_n(k)$ are the relators $\theta_*(P_{ij})$ with $j \le k$, thus the kernel is $\mathcal{J}_n(k)$ as in the statement of the proposition. □

Our "Birman kernel" $\mathcal{J}_n(k)$ is rather worse behaved than the analogous Birman kernel of Day and Putman. Their kernel, denoted $\mathcal{K}_{n,k,l}$, is finitely generated, whereas it may be shown by adapting the proof of their Theorem E that $\mathcal{J}_n(k)$ is not. This difference is due in part to the fact that their version of $\mathrm{P}\Pi\mathrm{A}_n(k)$ need only fix each of $x_1, \dots, x_k$ up to conjugacy. The lack of finite generation of $\mathcal{J}_n(k)$ is, however, not an obstacle to the goal of the current paper; 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 $v_1, \dots, v_k \in \mathbb{Z}^n$, and by $\mathcal{H}_n(k)$ the group $\mathrm{Hom}(\mathbb{Z}^{n-k}, (2\mathbb{Z})^k)$.

Lemma 3.5 *Fix $0 \le k \le n$. Then there exists the commutative diagram*

$$\begin{array}{ccccccccc}
1 & \longrightarrow & \mathcal{J}_n(k) & \longrightarrow & \mathrm{P}\Pi\mathrm{A}_n(k) & \underset{s}{\rightleftarrows} & \mathrm{P}\Pi\mathrm{A}_{n-k} & \longrightarrow & 1 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
1 & \longrightarrow & \mathcal{H}_n(k) & \longrightarrow & \Gamma_n[2](k) & \underset{t}{\rightleftarrows} & \Gamma_{n-k}[2] & \longrightarrow & 1
\end{array}$$

of split short exact sequences, where s and t are the obvious splitting homomorphisms.

Proof The top row is given by Proposition 3.4. A generating set for $\Gamma_n[2](k)$ follows from the proof of Lemma 2.4; 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 2.3. The bottom row then follows by an argument similar to the proof of Proposition 3.4, noting that the kernel is generated by the images of P_{ij} $(i > k,\ j \le k)$. It is straightforward to verify that this kernel is $\mathrm{Hom}(\mathbb{Z}^{n-k}, (2\mathbb{Z})^k)$. Intuitively, $\alpha \in \mathrm{Hom}(\mathbb{Z}^{n-k}, (2\mathbb{Z})^k)$ encodes how many (even) multiples of v_j $(1 \le i \le k)$ are added to each v_i $(k < j \le n)$.

The only vertical map left to consider is the right-most one. Its existence and surjectivity follow from Lemma 2.4. 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. □

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 $\alpha_k\colon \mathrm{P\Pi A}_n(k) \to \mathrm{Aut}(\mathcal{H}_n(k))$. Setting $k = 1$, we obtain the following lemma.

Lemma 3.6 *The group $\mathcal{J}_n(1) \cap \mathcal{PI}_n$ is normally generated in $\mathcal{J}_n(1)$ by the set*

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

Proof By Lemma 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 $Y := \{\phi P_{j1}\phi^{-1} \mid \phi \in \mathrm{P\Pi A}_n(1), 1 < j \le n\}$ generates $\mathcal{J}_n(1)$ by Proposition 3.4. Let a_j denote the image of P_{j1} in $\mathrm{GL}(n,\mathbb{Z})$. A direct calculation verifies that the set $\{a_j\}$ is 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

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

together with the generating set $\bar{Y}$, forms 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.

The group $\mathcal{J}_n(1) \cap \mathcal{PI}_n$ 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 2.3. It can be shown that the obvious lift of the finite set of relators

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

suffices to normally generate $\mathcal{J}_n(1) \cap \mathcal{PI}_n$. This may be seen using a simple induction argument on the length of a given expression of $\phi \in \mathrm{P}\Pi\mathrm{A}_n(1)$ on $C^{\pm 1}$.

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. □

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

Proof of Theorem A Recall that the set of doubled commutator transvections in $\mathcal{PI}_n$ is precisely the conjugacy class of $[P_{12}, P_{13}]$ in $\Pi\mathrm{A}_n$, and that the set of separating π–twists in $\mathcal{PI}_n$ is precisely the conjugacy class of

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

in $\Pi\mathrm{A}_n$.

The group $\mathcal{PI}_n$ acts on $\mathfrak{B}_n^\pi$ simplicially and without rotations. Combining Theorems 3.1, 3.2 and 3.3, we conclude that, for $n \geq 4$, $\mathcal{PI}_n$ is generated by the vertex stabilisers of the action on $\mathfrak{B}_n^\pi$.

Let $\mathcal{PI}_n(1)$ denote 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 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 5, we verify that the presentation of $\Gamma_3[2]$ given in Corollary 1.1 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 1–7 are trivial when lifted. Relator 8 lifts to $[P_{ij}, P_{ik}]$ and relator 9 lifts to $[P_{jk}, P_{ij}]{P_{ik}}^{-2}$, which equals $P_{ik}[P_{ij}, P_{ik}]{P_{ik}}^{-1}$. Thus the lifts of relators 8 and 9 are conjugate to $[P_{12}, P_{13}]$ in $\Pi\mathrm{A}_3$. 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, as each relator lifts to either a doubled commutator transvection, a separating π–twist or the identity automorphism.

Now suppose $n > 3$. By induction, the group $\mathcal{PI}_{n-1}$ is generated by the purported generating set. We lift this generating set to $\mathcal{PI}_n(1)$ in the obvious way.

By Lemma 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 the doubled commutator transvection $[P_{12}, P_{13}]$. For the latter, observe that

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

which again is conjugate in $\Pi\mathrm{A}_n$ to $[P_{12}, P_{13}]$, so we are done. □

Theorem A allows us to conclude that $\mathcal{PI}_n$ is normally generated in $\Pi\mathrm{A}_n$ by the automorphisms $\chi_1 = [P_{12}, P_{13}]$ and

$$\chi_2 = (P_{23}{P_{13}}^{-1}P_{31}P_{32}P_{12}{P_{21}}^{-1})^2.$$

Let $\Omega_n \leq \Pi\mathrm{A}_n$ denote the symmetric group on X. The presentation for $\Gamma_n[2] \cong \mathrm{P}\Pi\mathrm{A}_n/\mathcal{PI}_n$ given in Corollary 1.1 follows from Theorem A by adding the Ω_n–orbits of χ_1 and χ_2 to Collins' presentation for $\mathrm{P}\Pi\mathrm{A}_n$ as relators, then applying the obvious Tietze transformations.

We now demonstrate that the presence of separating π–twists in our generating set for $\mathcal{PI}_n$ is necessary.

Proposition 3.7 *For $n \geq 3$, the group generated by doubled commutator transvections is a proper subgroup of $\mathcal{PI}_n$.*

Proof Let $\mathcal{D}$ denote the subgroup of $\mathcal{PI}_n$ generated by doubled commutator transvections. In other words, $\mathcal{D}$ is the normal closure of $\chi_1 = [P_{12}, P_{13}]$ in $\Pi\mathrm{A}_n$. Then the Ω_n–orbit of χ_1 is a normal generating set for $\mathcal{D}$ in $\mathrm{P}\Pi\mathrm{A}_n$. Adding the members of this orbit to the presentation of $\mathrm{P}\Pi\mathrm{A}_n$ as relators yields a finite presentation $\mathcal{Q}$ of $\mathrm{P}\Pi\mathrm{A}_n/\mathcal{D}$, which may be altered using Tietze transformations so that it looks like the presentation in Corollary 1.1, with relator 10 (and relator 7, if $n = 3$) removed (where we interpret S_{ij} and O_i as formal symbols, rather than matrices). We shall show that the relations of $\mathcal{Q}$ are not a complete set of relations on the generating set $\{S_{ij}, O_i\}$ for $\Gamma_n[2] \cong \mathrm{P}\Pi\mathrm{A}_n/\mathcal{PI}_n$, and so conclude that $\mathcal{D} \neq \mathcal{PI}_n$.

It is easily shown that for

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

the image of χ_2 in $\Gamma_n[2]$, is trivial, but we shall show that ξ is non-trivial in the group presented by $\mathcal{Q}$. Observe that by trivialising all the generators of $\Gamma_n[2]$ except for S_{12} and S_{21}, we surject $\Gamma_n[2]$ 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 ξ under this map is $ABAB \neq 1$, and so ξ is non-trivial in the group presented by $\mathcal{Q}$. Therefore $\mathcal{D}$ is a proper subgroup of $\mathcal{PI}_n$. □

Note that in the proof of Proposition 3.7 we also showed that relators 1–9 of Corollary 1.1 are not a sufficient set of relators that hold between the O_i and S_{jk}, as relator 10 is not a consequence of the others. This allows us to conclude that the quotient space $\mathfrak{B}_3^\pi/\mathcal{PI}_3$ is not simply connected.

Corollary 3.8 *The complex $\mathfrak{B}_3^\pi/\mathcal{PI}_3$ is not simply connected.*

Proof By Theorem 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 of $\mathcal{PI}_3$ on $\mathfrak{B}_3^\pi$. As in the proof of Theorem A, the group generated by the vertex stabilisers of this action may be normally generated in $\Pi\mathrm{A}_3$ by the group $\mathcal{PI}_3(1)$. The same calculations as in the proof of Theorem A show that $\mathcal{PI}_3(1)$ is the normal closure of the doubled commutator transvection $[P_{12}, P_{13}]$. However, Proposition 3.7 showed that this normal closure is a proper subgroup of $\mathcal{PI}_3$, so the quotient $\mathfrak{B}_3^\pi/\mathcal{PI}_3$ is not simply connected. □

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 and Putman [11], while the latter is shown to be closely related to a complex already studied by Charney [7], which is $(n-3)$–connected.

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. Similarly, we also insist that the identity element of G is excluded from S.

We establish Theorem 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 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 $\Pi\mathrm{A}_n$ given in Section 2, to obtain the following list of defining relators for $\Pi\mathrm{A}_n$ on the generating set Z:

(1) $t_{ij} = t_{ji}$,
(2) ${t_{ij}}^2 = 1$,
(3) $u t_{ij} u^{-1} = t_{u(i)u(j)}$,
(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 both 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 $\Pi\mathrm{A}_n = \mathrm{E}\Pi\mathrm{A}_n \rtimes \Omega^{\pm 1}(X)$ given in Section 2. Note that relations 1–3 are a complete set of relations for the symmetric group, when generated by the transpositions $\{t_{ij}\}$ [25].

We now consider the Cayley graph $\mathrm{Cay}(\Pi\mathrm{A}_n, Z)$. Observe that for each $z \in Z$ 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 $\mathrm{Cay}(\Pi\mathrm{A}_n, Z)$ 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 $\Pi\mathrm{A}_n$ on $\mathrm{Cay}(\Pi\mathrm{A}_n, Z)$ and $\mathfrak{B}_n^\pi$, we can extend this map to a map of complexes $\Psi\colon \mathrm{Cay}(\Pi\mathrm{A}_n, Z) \to \mathfrak{B}_n^\pi$. Explicitly, Ψ takes a vertex $z_1 \cdots z_r$ of $\mathrm{Cay}(\Pi\mathrm{A}_n, Z)$ ($z_i \in Z^{\pm 1}$) to the vertex $z_1 \cdots z_r(x_1)$.

Proof of Theorem 3.1 This proof is modelled on Day and Putman's proof of [11, Theorem A]. Let

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

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

$$x_1 - z_1(x_1) - z_1z_2(x_1) - \cdots z_1z_2 \cdots - z_k(x_1) = x_1.$$

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 [19, Proposition 4.2]. We now verify that the loops in $\pi_1(\mathrm{Cay}(\Pi\mathrm{A}_n, Z), 1)$ corresponding to the relators in the above streamlined presentation for $\Pi\mathrm{A}_n$ 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 $P_{21}w = wP_{21}$ and $\iota_2 v = v\iota_2$. 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 $x_1x_2x_1$, 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 and Putman's proof [11] for the surjectivity of Ψ_*. We represent a loop $\omega \in \pi_1(\mathfrak{B}_n^\pi, x_1)$ as

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

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

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

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

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

By induction, to prove the claim all we need find is $\phi_k \in \Pi\mathrm{A}_n(1)$ such that

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

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

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

as required.

Now we define

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

so that

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

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. We are thus able to write

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

for some $z_j^i \in Z^{\pm 1}$ $(1 \le i \le k+1,\ 1 \le j \le p_i)$, 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

$$x_1 - \phi_1 t_{12}(x_1) - \cdots - \phi_1 t_{12} \cdots \phi_k t_{12}(x_1) = x_1,$$

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

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 $\{u_1, \dots, u_k\}$ of free abelian bases of $\mathbb{Z}^n$. Writing members of $\mathbb{Z}^n$ multiplicatively, there is an analogous notion of an odd 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 $\mathfrak{B}_n^\pi(\mathbb{Z})$. Just as ΠA_n acts transitively on the set of π–bases of F_n, so does $\Gamma_n[2]$ act transitively on the set of π–bases of $\mathbb{Z}^n$, as we now verify.

Lemma 4.1 *The group $\Gamma_n[2]$ acts transitively on the set of π–bases of $\mathbb{Z}^n$.*

Proof By definition, any π–basis is of the form $\{Mv_1, \dots, Mv_n\}$, for $M \in \Gamma_n[2]$ and $\{v_1, \dots, v_n\}$ the standard basis of $\mathbb{Z}^n$, where v_i has 1 in the i^{th} position and 0s elsewhere. Thus, we have a well-defined action of $\Gamma_n[2]$ on the set of π–bases of $\mathbb{Z}^n$ by left-multiplication of basis elements, which is transitive, as every π–basis lies in the same orbit as $\{v_1, \dots, v_n\}$. □

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. To prove the former, the following lemma is required.

Lemma 4.2 *Fix $\{u_1, \dots, u_n\}$ as a π–basis for $\mathbb{Z}^n$, and let $\rho\colon F_n \to \mathbb{Z}^n$ be the abelianisation map. Let $\widetilde{U} = \{\tilde{u}_1, \dots, \tilde{u}_k\}$ be a partial π–basis of F_n such that $\rho(\tilde{u}_i) = u_i$ for each $1 \le i \le k$. Then we can extend $\widetilde{U}$ to a π–basis of F_n, $\{\tilde{u}_1, \dots, \tilde{u}_n\}$, such that $\rho(\tilde{u}_i) = u_i$ for $1 \le i \le n$.*

Proof Extend $\{\tilde{u}_1, \ldots, \tilde{u}_k\}$ to a full π–basis of F_n, $\{\tilde{u}_1, \ldots, \tilde{u}_k, \tilde{u}'_{k+1}, \ldots, \tilde{u}'_n\}$, and define $u'_j = \rho(\tilde{u}'_j)$ for $k+1 \leq j \leq n$. Then $\{u_1, \ldots, u_k, u'_{k+1}, \ldots, u'_n\}$ is a π–basis for $\mathbb{Z}^n$. By Lemma 4.1, 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 $\phi(u'_j) = u_j$ for $k+1 \leq j \leq n$. By Lemma 3.5, ϕ lifts to some $\tilde{\phi} \in \mathrm{P\Pi A}_n(k)$, and the π–basis $\{\tilde{u}_1, \ldots, \tilde{u}_k, \tilde{\phi}(\tilde{u}'_{k+1}), \ldots, \tilde{\phi}(\tilde{u}'_n)\}$ projects onto $\{u_1, \ldots, u_n\}$ as desired. □

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

Theorem 4.3 *The spaces $\mathfrak{B}_n^\pi / \mathcal{PI}_n$ and $\mathfrak{B}_n^\pi(\mathbb{Z})$ are isomorphic as simplicial complexes.*

Proof Let $\rho\colon F_n \to \mathbb{Z}^n$ be the abelianisation map, and define a map of simplicial complexes $\Phi\colon \mathfrak{B}_n^\pi \to \mathfrak{B}_n^\pi(\mathbb{Z})$ 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 4.2, each π–basis of $\mathbb{Z}^n$ is the image of 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 4.2, 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. □

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 3.2 By a *unimodular sequence* in $\mathbb{Z}^n$, we mean an (ordered) sequence $(u_1, \ldots, u_k) \subset (\mathbb{Z}^n)^k$ 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 considers (among others) the subposet of sequences $(u_1, \ldots, u_k)$ such that each u_i is congruent to a standard basis vector v_j under mod 2 reduction of the entries of u_i. We denote by $\mathcal{X}_n$ the poset complex given by the subposet of such sequences. Theorem 2.5 of Charney [7] 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 $h\colon \mathcal{X}_n \to \mathfrak{B}_n^\pi(\mathbb{Z})^*$ by $(u_1,\dots,u_k) \mapsto \{u_1,\dots,u_k\}$. 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, as there might be some vertex $(u_1,\dots,u_k)$ of $\mathcal{X}_n$ such that $\{u_1,\dots,u_k\}$ 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, since $\pi_i(\mathcal{X}_n) = 1$ for $0 \leq i \leq n-3$. □

5 A presentation for $\Gamma_3[2]$

In order to apply Armstrong's theorem [1], it must be the case that $\mathfrak{B}_n^\pi/\mathcal{PI}_n \cong \mathfrak{B}_n^\pi(\mathbb{Z})$ is simply connected. However, as we have seen from Corollary 3.8, the space $\mathfrak{B}_3^\pi(\mathbb{Z})$ has non-trivial fundamental group. The case $n=3$ forms the base case of our inductive proof of Theorem A, 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 $\mathcal{PI}_3$.

The augmented partial π–basis complex for $\mathbb{Z}^3$ By adding simplices to the complex $\mathfrak{B}_3^\pi(\mathbb{Z})$, we obtain a simply connected complex that $\Gamma_3[2]$ acts on. This action allows us to present $\Gamma_3[2]$.

Recall that $\mathcal{B}_n(\mathbb{Z})$ is the partial basis complex of $\mathbb{Z}^n$. We represent its vertices by column vectors $u = (u^{(1)},\dots,u^{(n)})^T$. For use in the proof of Theorem 5.1, we follow Day and Putman [11] 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 5.1 *The complex $\mathcal{Y}$ is simply connected.*

Proof By Theorem 2.5 of Charney [7], 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 and Putman [11].

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 the group generated by $\Gamma_3[2]$ and the matrix

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

acts simplicially on $\mathcal{Y}$ and transitively on the 0–skeleton of $\mathcal{Y}$. This action is transitive on vertices because $\Gamma_3[2]$ acts transitively on the vertices of $\mathfrak{B}_3^{\pi}(\mathbb{Z})$, and any vertex of $\mathcal{Y} \setminus \mathfrak{B}_3^{\pi}(\mathbb{Z})$ may be taken to a vertex of $\mathfrak{B}_3^{\pi}(\mathbb{Z})$ by acting on it with E.

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 of $\mathcal{Y}$). 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 $(u^{(1)}, u^{(2)}, u^{(3)})^T$ to be a member of $\mathrm{lk}_{\mathcal{Y}}(v_3)$ is that $(u^{(1)}, u^{(2)})^T$ is a vertex of $\mathcal{B}_2(\mathbb{Z})$. The link $\mathrm{lk}_{\mathcal{Y}}(v_3)$ may thus be described as follows: it has one vertex for each pair (a,b), where a is a vertex of $\mathcal{B}_2(\mathbb{Z})$ and $b \in 2\mathbb{Z}$, with vertices (a,b) and (c,d) joined by an edge if and only if a and c are joined by an edge in $\mathcal{B}_2(\mathbb{Z})$. 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

$$w_0 - w_1 - \cdots - w_r,$$

where w_i $(1 \le i \le r)$ are vertices of $\mathcal{Y}$, and $w_0 = w_r = v_3$. Our goal is to systematically homotope this loop so that the rank of each vertex in the sequence is 0. Such a loop may be contracted to the vertex v_3, and so is trivial in $\pi_1(\mathcal{Y})$.

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

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

in $\mathrm{lk}_{\mathcal{Y}}(w_i)$, as seen in Figure 7. Fix attention on some q_j $(1 \le j \le s)$. By the division algorithm, there exist $a_j, b_j \in \mathbb{Z}$ such that $R(q_j) = a_j \cdot R(w_i) + b_j$, with $0 \le b_j < R(w_i)$. As in the proof of Lemma 2.4, we wish to ensure that a_j is even, if possible. In all but

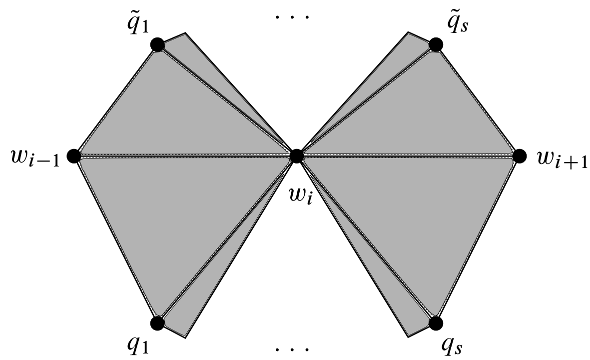

Figure 7: We find two homotopic paths that bound a disk inside $\mathrm{lk}_{\mathcal{Y}}(w_i)$, where the "upper" path seen here is constructed so that $R(\tilde{q}_j) < R(q_j)$ for $1 \le j \le s$.

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 $0 \le |B_j| < R(w_i)$. We do a case-by-case parity analysis. Note that since q_j and w_i are joined by an edge, $R(q_j)$ and $R(w_i)$ cannot both be odd, otherwise q_j and w_i would both map to the same member of $(\mathbb{Z}/2)^3$ when we reduce their entries mod 2. This would prohibit $\{q_j, w_i\}$ from extending to a basis J of $\mathbb{Z}^3$, otherwise the image of J in $(\mathbb{Z}/2)^3$ would generate despite only having two members. 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 $B_j = b_j - R(w_i)$. In that case, $|B_j| < R(w_i)$, since b_j must be odd and hence non-zero. 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

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

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

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

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

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

as seen in Figure 7. By construction, $R(\tilde{q}_j) < R(w_i)$. Iterating this procedure continually homotopes ω until it is inside the contractible (full) subcomplex spanned by v_3 and $\mathrm{lk}_{\mathcal{Y}}(v_3)$, and hence is trivial. Therefore $\pi_1(\mathcal{Y}) = 1$. □

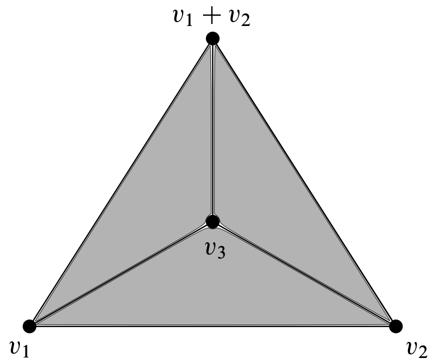

Figure 8: The quotient complex of $\mathcal{Y}$ under the action of $\Gamma_3[2]$. We have labelled its vertices using representatives from the vertex set of $\mathcal{Y}$.

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 Corollary 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 $b_j = 0$, as $\tilde{q}_j \notin \mathfrak{B}_3^\pi(\mathbb{Z})$. 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.

Presenting $\Gamma_3[2]$ Let $\Gamma_3[2](w_1, \dots, w_k)$ denote the stabiliser of the ordered tuple $(w_1, \dots, w_k)$ of vertices of $\mathcal{Y}$. 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 8. The quotient lifts to a subcomplex W of $\mathcal{Y}$ via the vertex labels seen in Figure 8. This subcomplex is what Brown [6] refers to as a *fundamental domain* for the action, and so a theorem of Brown [6, Theorem 3] allows us to conclude that $\Gamma_3[2]$ is the free product of the stabilisers of the vertices of W modulo *edge relations*, which identify the copies of the edge stabiliser $\Gamma_3[2](a, b)$ inside the vertex stabilisers $\Gamma_3[2](a)$ and $\Gamma_3[2](b)$, where $a, b \in \{v_1, v_2, v_3, v_1 + v_2\}$ are distinct.

We obtain a finite presentation for $\Gamma_3[2](v_1)$ using the semi-direct production decomposition of $\Gamma_3[2](v_1)$ given by Lemma 3.5 (noting that $\Gamma_2[2] \cong \mathrm{P\Pi A}_2$). 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 all relators of the form 1–9 (excluding 7, as it is not defined when $n = 3$) seen in Corollary 1.1. By permuting the indices accordingly, we also obtain finite presentations for the stabiliser groups $\Gamma_3[2](v_2)$ and $\Gamma_3[2](v_3)$. Identifying the edge stabiliser subgroups of these three groups appropriately, we obtain the presentation seen in Corollary 1.1 without relators 7 and 10; we denote this presentation by $\mathcal{P}$.

We now see that the effect of identifying the edge stabiliser subgroups of $\Gamma_n[2](v_1+v_2)$ with the corresponding copies inside the other three vertex stabiliser groups is to include one additional relator: relator 10. Since $\Gamma_3[2](v_1+v_2)$ and $\Gamma_3[2](v_1)$ are conjugate inside $\mathrm{GL}(3,\mathbb{Z})$, we take a formal presentation for $\Gamma_3[2](v_1+v_2)$ by adding a "hat" to each of the symbols in the presentation of $\Gamma_3[2](v_1)$.

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 1 in the $(2,1)$ position. In Table 1 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,

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

Generator M of $\Gamma_3[2](v_1)$	The conjugate $\widehat{M} = E_{21}\cdot M\cdot {E_{21}}^{-1}$
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 1: The conjugates of the generating set of $\Gamma_3[2](v_1)$ by E_{21}

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 the relations corresponding to this edge stabiliser, we must express them using the generators of $\Gamma_3[2](v_1+v_2)$, and set them to be equal accordingly. Consulting Table 1, we get the edge relations

$$\widehat{O}_3 = O_3,\quad \widehat{S}_{13} = S_{13}S_{23}\quad\text{and}\quad \widehat{S}_{23} = S_{23}.$$

Note that these relations simply reiterate the expressions we had already determined for $\widehat{O}_3$, $\widehat{S}_{13}$ and $\widehat{S}_{23}$. 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

$$\widehat{O}_2 = S_{21}O_2,\quad \widehat{S}_{12} = O_1O_2S_{21}{S_{12}}^{-1}\quad\text{and}\quad \widehat{S}_{32} = S_{32}{S_{31}}^{-1}.$$

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 $\widehat{O}_3 = O_3$, $S_{13} = \widehat{S}_{13}\widehat{S}_{23}^{-1}$ and $\widehat{S}_{13} = S_{13}S_{23}$. 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.

We now use these edge relations to replace the formal relators defining $\Gamma_3[2](v_1 + v_2)$ with words on the generating set $\{S_{ij}, O_k\}$. Using Tietze transformations and Brown's Theorem 3 [6], we may then conclude that a complete presentation for $\Gamma_3[2]$ is obtained by adding these relators to the presentation $\mathcal{P}$. For example, the relator $\widehat{O}_2^2$ becomes $(S_{21}O_2)^2$. All but one of these additional relators are consequences of ones already in $\mathcal{P}$. The one relator that is not is $[\widehat{S}_{13}, \widehat{S}_{32}]\widehat{S}_{12}^{-2}$, which becomes

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

Using the other relations in $\Gamma_3[2]$, this word may be rewritten in the form of relator 10 in Corollary 1.1; we have thus verified that the presentation given in Corollary 1.1 is correct when $n = 3$. This proves Proposition 1.2.

References

[1] **M A Armstrong**, *On the fundamental group of an orbit space*, Proc. Cambridge Philos. Soc. 61 (1965) 639–646 MR0187244

[2] **J S Birman**, *Braids, links, and mapping class groups*, Annals of Mathematics Studies 82, Princeton Univ. Press (1974) MR0375281

[3] **J S Birman**, **H M Hilden**, *On isotopies of homeomorphisms of Riemann surfaces*, Ann. of Math. 97 (1973) 424–439 MR0325959

[4] **T E Brendle**, **D Margalit**, *Point pushing, homology, and the hyperelliptic involution*, Michigan Math. J. 62 (2013) 451–473 MR3102525

[5] **T Brendle**, **D Margalit**, **A Putman**, *Generators for the hyperelliptic Torelli group and the kernel of the Burau representation at $t = -1$*, preprint (2014) `arXiv:1211.4018`

[6] **K S Brown**, *Presentations for groups acting on simply-connected complexes*, J. Pure Appl. Algebra 32 (1984) 1–10 MR739633

[7] **R Charney**, *On the problem of homology stability for congruence subgroups*, Comm. Algebra 12 (1984) 2081–2123 MR747219

[8] **D J Collins**, *Palindromic automorphisms of free groups*, from: "Combinatorial and geometric group theory", (A J Duncan, N D Gilbert, J Howie, editors), London Math. Soc. Lecture Note Ser. 204, Cambridge Univ. Press (1995) 63–72 MR1320275

[9] **J Crisp**, **L Paris**, *Artin groups of type B and D*, Adv. Geom. 5 (2005) 607–636 MR2174484

[10] **M Day**, **A Putman**, *A Birman exact sequence for* $\mathrm{Aut}(F_n)$, Adv. Math. 231 (2012) 243–275 MR2935388

[11] **M Day**, **A Putman**, *The complex of partial bases for* F_n *and finite generation of the Torelli subgroup of* $\mathrm{Aut}(F_n)$, Geom. Dedicata 164 (2013) 139–153 MR3054621

[12] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Mathematical Series 49, Princeton Univ. Press (2012) MR2850125

[13] **N J Fullarton**, *Palindromic automorphisms of free groups and rigidity of automorphism groups of right-angled Artin groups*, PhD thesis, University of Glasgow (2014) MR3389544 Available at `http://search.proquest.com/docview/1689623221`

[14] **N J Fullarton**, **A Thomas**, *Palindromic automorphisms of right-angled Artin groups*, preprint (2015) `arXiv:1510.03939`

[15] **H H Glover**, **C A Jensen**, *Geometry for palindromic automorphism groups of free groups*, Comment. Math. Helv. 75 (2000) 644–667 MR1789180

[16] **R Hain**, *Finiteness and Torelli spaces*, from: "Problems on mapping class groups and related topics", (B Farb, editor), Proc. Sympos. Pure Math. 74, Amer. Math. Soc. (2006) 57–70 MR2264131

[17] **C Jensen**, **J McCammond**, **J Meier**, *The Euler characteristic of the Whitehead automorphism group of a free product*, Trans. Amer. Math. Soc. 359 (2007) 2577–2595 MR2286046

[18] **R Kobayashi**, *A finite presentation of the level* 2 *principal congruence subgroup of* $\mathrm{GL}(n;\mathbb{Z})$, preprint (2015) `arXiv:1407.1612`

[19] **R C Lyndon**, **P E Schupp**, *Combinatorial group theory*, Ergeb. Math. Grenzgeb. 89, Springer, Berlin (1977) MR0577064

[20] **W Magnus**, **A Karrass**, **D Solitar**, *Combinatorial group theory*, 2nd edition, Dover, New York (1976) MR0422434

[21] **J D McCarthy**, **U Pinkall**, *Representing homology automorphisms of nonorientable surfaces* (2004) Revised version of Max Planck Inst. preprint MPI/SFB 85-11 Available at `http://users.math.msu.edu/users/mccarthy/publications/repr.hom.auts.pdf`

[22] **J Milnor**, *Introduction to algebraic K–theory*, Annals of Mathematics Studies 72, Princeton Univ. Press (1971) MR0349811

[23] **T Morifuji**, *On Meyer's function of hyperelliptic mapping class groups*, J. Math. Soc. Japan 55 (2003) 117–129 MR1939188

[24] **B Perron**, **J P Vannier**, *Groupe de monodromie géométrique des singularités simples*, Math. Ann. 306 (1996) 231–245 MR1411346

[25] **A Putman**, *Obtaining presentations from group actions without making choices*, Algebr. Geom. Topol. 11 (2011) 1737–1766 MR2821439

[26] **R D Wade**, *Folding free-group automorphisms*, Q. J. Math. 65 (2014) 291–304 MR3179663

Department of Mathematics, Rice University
MS 136, 6100 Main Street, Houston, TX 77005, USA

neil.fullarton@rice.edu

Received: 4 November 2014 Revised: 14 April 2015