THE PURE SYMMETRIC AUTOMORPHISMS OF A FREE GROUP FORM A DUALITY GROUP

NOEL BRADY, JON McCAMMOND, JOHN MEIER, AND ANDY MILLER

Abstract. The pure symmetric automorphism group of a finitely generated free group consists of those automorphisms which send each standard generator to a conjugate of itself. We prove that these groups are duality groups.

1. Introduction

Let F_n be a finite rank free group with fixed free basis $X = \{x_1, \ldots, x_n\}$. The *symmetric automorphism group* of F_n, hereafter denoted Σ_n, consists of those automorphisms that send each $x_i \in X$ to a conjugate of some $x_j \in X$. The *pure symmetric automorphism group*, denoted $\mathrm{P}\Sigma_n$, is the index $n!$ subgroup of Σ_n of symmetric automorphisms that send each $x_i \in X$ to a conjugate of itself. The quotient of $\mathrm{P}\Sigma_n$ by the inner automorphisms of F_n will be denoted $\mathrm{OP}\Sigma_n$. In this note we prove:

Theorem 1.1. *The group* $\mathrm{OP}\Sigma_n$ *is a duality group of dimension* $n-2$*.*

Corollary 1.2. *The group* $\mathrm{P}\Sigma_n$ *is a duality group of dimension* $n-1$*, hence* Σ_n *is a virtual duality group of dimension* $n-1$*.*

(In fact we establish slightly more: the dualizing module in both cases is $\mathbb{Z}$-free.)

Corollary 1.2 follows immediately from Theorem 1.1 since F_n is a 1-dimensional duality group, there is a short exact sequence

$$1 \to F_n \to \mathrm{P}\Sigma_n \to \mathrm{OP}\Sigma_n \to 1$$

and any duality-by-duality group is a duality group whose dimension is the sum of the dimensions of its constituents (see Theorem 9.10 in [2]).

That the virtual cohomological dimension of Σ_n is $n-1$ was previously established by Collins in [9]. In more recent work, Gutiérrez and Krstić have shown that Σ_n has a regular language of normal forms [11], and Orlandi-Korner has computed the BNS-invariant of $\mathrm{P}\Sigma_n$ [13]. Bogley and Krstić have recently announced a computation of the integral homology of the groups Σ_n [4].

We establish Theorem 1.1 using a variation on the Main Theorem of [8] which is presented in Section 3. To apply this type of theorem, we need a contractible space on which $\mathrm{OP}\Sigma_n$ acts where the cell stabilizers are duality groups of the appropriate dimensions. Such a space has, in fact, already been constructed by McCullough and Miller in [12]. In Section 4 we review their construction and the properties of the stabilizers. Lastly we need to establish that the fundamental domain for this action is a Cohen-Macaulay complex. This is shown in Section 5. The final section contains various open questions related to these groups.

Date: April 3, 2001.

We thank Laura Anderson, Ken Brown and Craig Jensen for helpful discussions during the development of this paper.

2. Duality Groups

In [3] Bieri and Eckmann introduced a class of groups, called duality groups, whose cohomology behaves like the cohomology of compact manifolds.

Definition 2.1 (Duality groups)**.** Let G be an FP group of cohomological dimension n. The group G is an *n-dimensional duality group* if there exists a G-module D (called the dualizing module) such that $H^i(G, M) \simeq H_{n-i}(G, D \otimes M)$ for all integers i and for all G-modules M. Equivalently, G is a duality group if its cohomology with group ring coefficients is torsion free and concentrated in dimension n. The dualizing module is this case is $H^n(G, \mathbb{Z}G)$. A group G is said to be a *virtual* duality group if it has a finite index subgroup that is a duality group. See [2] and [7] for further information on duality properties for groups.

Example 2.2. The simplest examples of duality groups are the free and free abelian groups. Finitely generated free groups are 1-dimensional duality groups and the free abelian group $\mathbb{Z}^n$ is an n-dimensional duality group since it admits a manifold $K(\pi, 1)$ of dimension $n = \text{cd}(\mathbb{Z}^n)$. The braid group B_n is a duality group of dimension $n - 1$ (this is discussed in section 6), and the group $\text{Aut}(F_n)$ is a virtual duality group of dimension $2n - 2$ [1]. Since the braid group B_n is a subgroup of Σ_n, and Σ_n is a subgroup of $\text{Aut}(F_n)$, Σ_n is sandwiched between two virtual duality groups, thus lending credence to our Main Theorem.

3. Actions on posets

In this paper we use a theorem on group actions on simplicial complexes arising from actions on posets that was developed (but not stated) in [8]. Before discussing this theorem we establish some terminology.

Definition 3.1 (Posets)**.** We will assume that all posets are finite dimensional, that is, that the geometric realizations $|\mathcal{P}|$ are finite dimensional. A poset $\mathcal{P}$ is *graded* if all its maximal chains have the same length. If ς is an element of a poset $\mathcal{P}$, the *rank* of ς is the length of an unrefinable chain from a minimal element of $\mathcal{P}$ to ς. Thus minimal elements have rank zero, and if $\mathcal{P}$ is graded, all maximal elements have rank equal to the dimension of $|\mathcal{P}|$. If $\mathcal{P}$ is graded, then one can define the *corank* of $\varsigma \in \mathcal{P}$ to be $\text{crk}(\varsigma) \equiv d - \text{rk}(\varsigma)$ where d is the dimension of $|\mathcal{P}|$. Although this definition of rank is slightly non-standard, it simplifies the notation in this context.

Definition 3.2 (Cohen-Macaulay)**.** A finite dimensional simplicial complex $\mathcal{K}$ is *Cohen-Macaulay* if

$$\widetilde{H}_i(lk(\sigma), \mathbb{Z}) = 0$$

for all simplices σ and for all $i < \dim(lk(\sigma))$. Note: We allow σ to be the empty simplex, in which case $lk(\sigma)$ is all of $\mathcal{K}$. A poset $\mathcal{P}$ is said to be *Cohen-Macaulay* if its geometric realization $|\mathcal{P}|$ is Cohen-Macaulay.

It's known that if G is a group of type FP, with $cd(G) = n$, and G acts on a contractible complex $\mathcal{X}$ where the stabilizer of each cell σ is an $(n - |\sigma|)$-dimensional duality group, then G is an n-dimensional duality group (Corollary to Theorem A in [8]). However, one often constructs group actions by having a group act on a poset, thereby getting an action on the geometric realization of that poset. In such a situation, the nice duality between the dimension of a cell and the dimension of its stabilizer will not hold. For example, each and every element of the poset will contribute a vertex to the geometric realization, thus the collection of vertex stabilizers will consist of the stabilizers of each and every element of the poset.

Recall that an action of G on a cellular complex $\mathcal{X}$ has a *strong fundamental domain* $\mathcal{F}$ if $\mathcal{F}$ is a subcomplex of $\mathcal{X}$ and the natural map $\mathcal{F} \to G\backslash\mathcal{X}$ is a bijection. We say that a G-poset $\mathcal{P}$ has a *strong fundamental domain* if there is a subposet $\mathcal{F} \subset \mathcal{P}$ which is a filter (if $\varsigma \in \mathcal{F}$ and $\tau > \varsigma$, then $\tau \in \mathcal{F}$) and which contains unique representatives of each G-orbit in $\mathcal{P}$. If $\mathcal{F}$ is a strong fundamental domain for the G-action on $\mathcal{P}$, it follows that $|\mathcal{F}|$ is a strong fundamental domain for the G-action on $|\mathcal{P}|$.

Theorem 3.3. *Let G be a group of type FP, with $cd(G) = n$. Let G act on a poset $\mathcal{P}$, whose geometric realization $|\mathcal{P}|$ is contractible, and where there is a strong fundamental domain $\mathcal{F} \subset \mathcal{P}$ that is finite and Cohen-Macaulay. If the stabilizer of each element $\varsigma \in \mathcal{P}$ is an $(n - crk(\varsigma))$-dimensional duality group, then G is an n-dimensional duality group.*

If in addition to the conditions above, it's also true that the dualizing modules of the stabilizers are all $\mathbb{Z}$-free, then the dualizing module of G is $\mathbb{Z}$-free.

The proof of this theorem is essentially given in §7 of [8], but there it was developed for a specific application to Artin groups, and it is presented in that context. A reader interested in this theorem, but not particularly interested in Artin groups, would find it difficult to generalize the argument in [8] in order to construct the proof of Theorem 3.3. Thus we quickly run through an outline of the proof.

Outline of proof. Because G is FP and $cd(G) = n$, in order to establish duality it suffices to establish that the cohomology of G with $\mathbb{Z}G$ coefficients is concentrated in dimension n and is $\mathbb{Z}$-torsion free.

We express $H^*(G, \mathbb{Z}G)$ in terms of the equivariant cohomology for the action of G on $|\mathcal{P}|$, $H^*_G(|\mathcal{P}|)$. The standard equivariant spectral sequence arises from filtering a space by skeleta (see for example §VII.7 in [7]). However, our filtration of $|\mathcal{P}|$ is by G-equivariant subcomplexes that are more naturally related to the underlying poset structure. We let $|\mathcal{P}|^0$ denote the subcomplex constructed using only corank 0 elements; this should be thought of as the collection of 'vertices' for this complex. In general, $|\mathcal{P}|^i$ consists of the geometric realization of the subposet consisting of corank k elements for $0 \leq k \leq i$. In essence, we are replacing the notion of "dimension" with the notion of "corank".

If $\varsigma \in \mathcal{P}$ then let $\mathcal{P}_{>\varsigma}$ and $\mathcal{P}_{\geq\varsigma}$ denote the subposets of all elements greater than (greater than or equal to) ς in $\mathcal{P}$. We let the stabilizer of ς under the action of G be denoted G_ς.

Because $\mathcal{F}$ is a strong fundamental domain, the relative chains $C(|\mathcal{P}|^p, |\mathcal{P}|^{p-1})$ can be expressed in terms of induced modules based at the fundamental domain:

$$\bigoplus_{\mathrm{crk}(\varsigma)=p, \varsigma\in\mathcal{F}} C\left(|\mathcal{F}_{\geq\varsigma}|, |\mathcal{F}_{>\varsigma}|\right) \uparrow_{G_\varsigma}^{G} .$$

Let F be a finite projective resolution of $\mathbb{Z}$ as a $\mathbb{Z}G$-module. Then

$$\begin{aligned}
&\mathrm{Hom}_G\left(F, \mathrm{Hom}_{\mathbb{Z}}(C(|\mathcal{P}|^p, |\mathcal{P}|^{p-1}), \mathbb{Z}G)\right)\\
&= \bigoplus_{\mathrm{crk}(\varsigma)=p, \varsigma\in\mathcal{F}} \mathrm{Hom}_G\left(F, \mathrm{Hom}_{\mathbb{Z}}\left(C\left(|\mathcal{F}_{\geq\varsigma}|, |\mathcal{F}_{>\varsigma}|\right)\uparrow^G_{G_\varsigma}, \mathbb{Z}G\right)\right)\\
&= \bigoplus_{\mathrm{crk}(\varsigma)=p, \varsigma\in\mathcal{F}} \mathrm{Hom}_{G_\varsigma}\left(F, \mathrm{Hom}_{\mathbb{Z}}\left(C\left(|\mathcal{F}_{\geq\varsigma}|, |\mathcal{F}_{>\varsigma}|\right), \mathbb{Z}G\right)\right)
\end{aligned}$$

It follows that

$$\begin{aligned}
E_1^{p,q} = H_G^{p+q}\left(\left(|\mathcal{P}|^p, |\mathcal{P}|^{p-1}\right), \mathbb{Z}G\right) &= \bigoplus_{\mathrm{crk}(\varsigma)=p, \varsigma\in\mathcal{F}} H_{G_\varsigma}^{p+q}\left(\left(|\mathcal{F}_{\geq\varsigma}|, |\mathcal{F}_{>\varsigma}|\right), \mathbb{Z}G\right)\\
&= \bigoplus_{\mathrm{crk}(\varsigma)=p, \varsigma\in\mathcal{F}} H^q(G_\varsigma, H^p\left(\left(|\mathcal{F}_{\geq\varsigma}|, |\mathcal{F}_{>\varsigma}|\right), \mathbb{Z}G\right))
\end{aligned}$$

Where the last equality uses the Cohen-Macaulay hypothesis. The pair $(|\mathcal{F}_{\geq\varsigma}|, |\mathcal{F}_{>\varsigma}|)$ is a cone over the link of the vertex associated with ς in $|\mathcal{F}_{\geq\varsigma}|$, relative to the link. If $\sigma = \varsigma_0 < \varsigma_1 < \cdots < \varsigma$ is an unrefinable chain in $\mathcal{F}$ starting with a minimal element and ending with the element ς, then the link of ς in $|\mathcal{F}_{\geq\varsigma}|$ is the same as the link of the simplex $|\sigma|$ associated with the chain σ in $|\mathcal{F}|$. The dimension of $|\sigma|$ is the rank of ς, thus by the Cohen-Macaulay hypothesis, this link has the homology of a wedge of spheres of dimension $d - \mathrm{rk}(\varsigma) - 1$ where d is the dimension of $|\mathcal{P}|$. So $(|\mathcal{F}_{\geq\varsigma}|, |\mathcal{F}_{>\varsigma}|)$ has the homology of a wedge of $(d - \mathrm{rk}(\varsigma))$-spheres. But $(d - \mathrm{rk}(\varsigma)) = \mathrm{crk}(\varsigma) = p$, hence the relative homology is trivial except in dimension p where it's free abelian. Thus

$$H^p((|\mathcal{F}_{\geq\varsigma}|, |\mathcal{F}_{>\sigma}|), \mathbb{Z}G) = \mathrm{Hom}(H_p(|\mathcal{F}|_{\geq\sigma}, |\mathcal{F}|_{>\sigma}), \mathbb{Z}G) = \bigoplus \mathbb{Z}G$$

Tracing back through the equalities we see

$$E_1^{p,q} = \prod_{\mathrm{crk}(\varsigma)=p, \varsigma\in\mathcal{F}} H^q(G_\varsigma, \bigoplus \mathbb{Z}G).$$

Because the stabilizer G_ς is an $(n-p)$-dimensional duality group, and $\mathbb{Z}G$ as a $\mathbb{Z}G_\varsigma$-module is a sum of $\mathbb{Z}G_\varsigma$'s, $E_1^{p,q}$ is $\mathbb{Z}$-torsion free when $q = n - p$, and is zero otherwise. It follows that the entire spectral sequence lies in total degree $p + q = n$. Because all the entries below the n^{th}-diagonal are zero, $H^q(G, \mathbb{Z}G) = 0$ below dimension n; since each G_ς is an FP group, cohomology commutes with direct sums, hence $H^n(G, \mathbb{Z}G)$ is $\mathbb{Z}$-torsion free. If the stabilizers all have $\mathbb{Z}$-free dualizing modules, then the $E_1^{p,q}$ term is actually $\mathbb{Z}$-free (for $p + q = n$), hence the dualizing module of G is also $\mathbb{Z}$-free. □

4. The McCullough-Miller Complex

In [12] McCullough and Miller introduced a family of contractible complexes which admit actions by certain automorphism groups of free products. They construct in particular a complex on which $\mathrm{OP}\Sigma_n$ acts cocompactly. The fundamental domain for the action of $\mathrm{OP}\Sigma_n$ on this space is the geometric realization of a finite poset associated to certain trees. This fundamental domain will be described first, followed by a description of the entire complex.

Definition 4.1 ([n]-labelled bipartite trees)**.** Let $[n] \equiv \{1, \ldots, n\}$. An *[n]-labelled bipartite tree* is a tree T together with a bijection from $[n]$ to a subset of its vertex set such that the image of $[n]$ includes all of the vertices of valence 1 and for every edge in T exactly one of its endpoints lies in the image of $[n]$. Several [4]-labelled bipartite trees are shown in Figure 1.

The vertices that lie in the image of $[n]$ are called *labelled vertices* and we use v_i to denote the vertex labelled by i. Two labelled bipartite trees are considered to be equivalent if there is a label preserving graph isomorphism between them. If there are m unlabelled vertices in a labelled bipartite tree T, then the *rank* of T is $m-1$. (Note: In [12] there is an alternative definition of "rank"; a short induction argument proves their definition is equivalent to our definition.)

Let T be a labelled bipartite tree with two distinct edges e_1 and e_2 which share a common labelled endpoint v and whose unlabelled endpoints are u_1 and u_2 respectively. Let T' be the tree obtained from T by identifying the edges e_1 and e_2 as well as the vertices u_1 and u_2. We say that the tree T' is obtained from T by *folding at v*. Notice that folding reduces the rank of a tree by 1. For example, in Figure 1, C can be folded at 2 to produce B, B can be folded at 3 to produce A, and D can be folded at 3 to produce B.

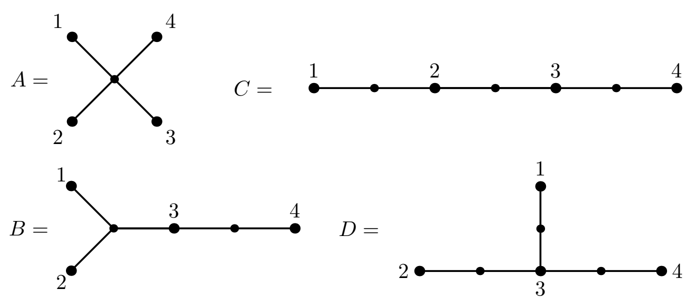

FIGURE 1. Examples of [4]-labelled bipartite trees.

Definition 4.2 (Whitehead poset)**.** The *Whitehead poset* W_n is the poset consisting of all [n]-labelled bipartite trees under the partial order induced by folding. Specifically, $T' \prec T$ if one can obtain T' from T by a sequence of foldings. In the poset W_n there is a single element of rank 0 called the *nuclear vertex*. It consists of an unlabelled vertex of valence n and n labelled vertices of valence 1. (See for example tree A in Figure 1.) Let T be an [n]-labelled bipartite tree. We note that the definitions of rank have been chosen so that the rank of T as an [n]-labelled bipartite tree (Definition 4.1) agrees with rank of T as an element of the poset W_n (Definition 3.1).

Lemma 4.3. *The Whitehead poset W_n is graded.*

Proof. Let T be an [n]-labelled bipartite tree. If T contains a labelled vertex v_i with valence greater than 1 then T can be folded at v_i to obtain a tree of smaller rank. This shows that repeatedly folding any such tree will eventually result in the nuclear element. In particular, every unrefinable chain must start at the nuclear element. On the other hand, if T contains an unlabelled vertex of valence greater than 2, then T can be unfolded to a tree of higher rank. Thus the other end of an unrefinable chain is the barycentric subdivision of a tree with exactly n vertices and $n-1$ edges. Since the unique minimum element in W_n has rank 0, all of the maximal elements have rank $n-1$, and every unrefinable chain changes the rank by 1 each time, every maximal chain has the same length. □

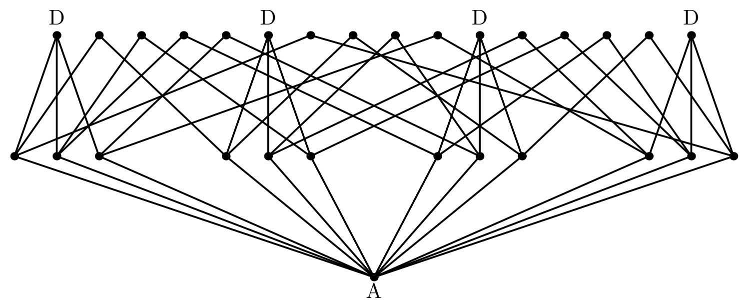

FIGURE 2. The poset W_4

Example 4.4 (W_4 and W_5)**.** The number of elements in W_n grows quite rapidly. For example, $|W_3| = 4$, $|W_4| = 29$ and $|W_5| = 311$. The poset W_4 is illustrated in Figure 2. The letters used to label an element of W_4 in Figure 2 are meant to indicate which of the trees in Figure 1 this element resembles when the [4]-labelling is ignored. Thus the four vertices labelled D represent trees isomorphic to the tree labelled D in Figure 1 and the vertex labelled A, the nuclear vertex for W_4, represents a tree isomorphic to the tree labelled A. Using this convention, all of the unlabelled vertices in the top row should be labelled C and all of the unlabelled vertices in the middle row should be labelled B. (See also Figure 8 in [12].)

The poset W_5 is too large to represent in the manner of Figure 2. However, the link of the nuclear vertex is a 2-complex with a natural label permuting action by the 5-element symmetric group. The fundamental domain under this action consists of 11 2-simplices as illustrated in Figure 3; the quotient of the link under this action is formed by identifying the three edges labelled with orientations. The rank of each vertex is indicated by its shape.

Definition 4.5 (Markings)**.** A *marking* of a labelled bipartite tree T consists of a basis of F_n, which we'll denote $\{y_1, \dots, y_n\}$, where the element y_i is a conjugate of $x_i \in X$, and is associated with the vertex labelled i in T.

Definition 4.6 (Marked automorphisms)**.** An automorphism $\alpha \in \mathrm{P}\Sigma_n$ is *carried* by a marked, $[n]$-labelled, bipartite tree T if:

1. There is an element y_i marking a vertex $v_i \in T$, and $\alpha(y_i) = y_i$;
2. For each vertex v_j $(j \neq i)$, $\alpha(y_j)$ is a conjugate of y_j via some power of y_i;
3. If v_j and v_k are in the same component of $T \backslash \{v_i\}$ then y_j and y_k are conjugated by the same power of y_i.

Definition 4.7 (MM_n)**.** The McCullough-Miller complex MM_n is the simplicial realization of the poset of marked, $[n]$-labelled, bipartite trees, modulo the equivalence relation generated by identifying two such trees if there is an automorphism carried by one that results in the other.

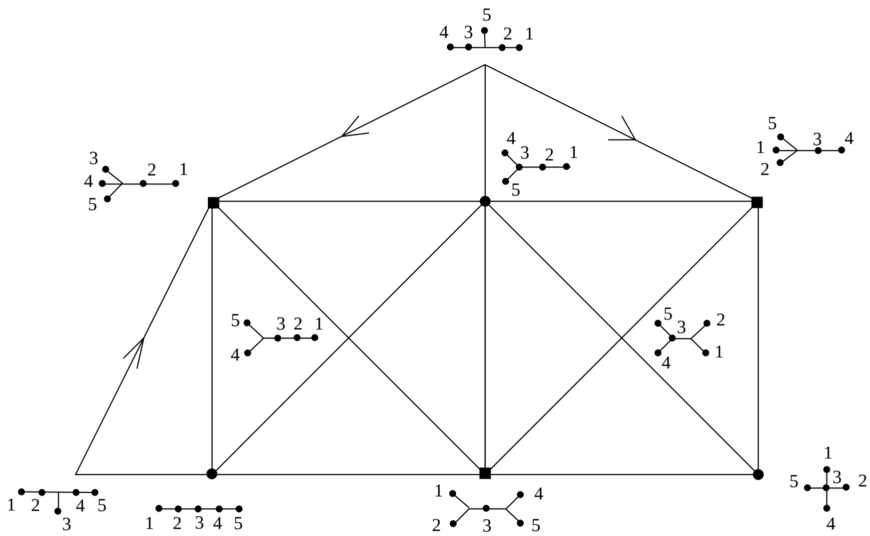

FIGURE 3. The quotient of the link of the nuclear vertex in W_5, under the action of the 5-element symmetric group.

Theorem 4.8 (McCullough-Miller [12])**.** *The complex MM$_n$ is a contractible complex of dimension $n-2$.*

The group PΣ_n acts on MM$_n$ by permuting the markings. The fundamental domain consists of the copy of W_n obtained by restricting to markings using the basis $X = \{x_1, \ldots, x_n\}$; it is a strong fundamental domain under the action of PΣ_n. Given the equivalence relation, the group of inner automorphisms acts trivially on MM$_n$, hence the action of PΣ_n yields an action of the quotient group OPΣ_n. The stabilizer of a vertex in MM$_n$, corresponding to a marked labelled bipartite tree T, consists of all automorphisms that can be expressed as a product of automorphisms that are carried by T. In Section 5 of [12], McCullough and Miller compute these stabilizers for various group actions. In the specific case of OPΣ_n, they show the following:

Lemma 4.9. *Under the action of* OPΣ_n*, the stabilizer of a rank k vertex of MM$_n$ is a free abelian group of rank k. In particular, this stabilizer is a k-dimensional duality group.*

5. The "Cohen-Macaulayness" of W_n

In order to apply Theorem 3.3, it only remains to establish that the poset W_n is Cohen-Macaulay. We do this via a technique of Björner and Wachs [6]. Specifically, we will establish that a closely related poset has a recursive atom ordering. Once this has been shown, the shellability and the "Cohen-Macaulayness" of W_n will follow easily using only standard results. First we recall some definitions.

Definition 5.1 (Shellable)**.** Let P be a finite, graded poset, and let $|P|$ be its geometric realization. The maximal simplices of $|P|$ correspond to maximal chains in P and are called *facets*. The poset P is *shellable* if the facets can be totally ordered $\{\sigma_1, \sigma_2, \ldots, \sigma_n\}$ so that $\sigma_{i+1} \cap \{\cup_{n=1}^{i} \sigma_n\}$ is a union of codimension one faces of σ_{i+1}.

It is well-known that shellability implies Cohen-Macaulay, so it is sufficient to establish that W_n is shellable. Actually, shellability implies a slightly stronger condition. Shellable simplicial complexes are *homotopy* Cohen-Macaulay, that is, $\pi_i(lk(\sigma))$ is trivial for all σ and all $i < dim(lk(\sigma))$, but this stronger statement is not necessary to establish duality. To show shellability we use recursive atom orderings.

Definition 5.2 (Recursive atom ordering)**.** In a poset, y *covers* x if $x < y$ and $x < z \leq y$ implies $z = y$. The *atoms* of a bounded, graded poset P are those elements that cover the minimum element $\hat{0}$. A bounded, graded poset P admits a *recursive atom ordering* if the length of P is one or if the length of P is greater than one and there is an ordering $a_1, \ldots, a_t$ of the atoms of P that satisfies:

1. For all $j = 1, 2, \ldots, t$, the interval $[a_j, \hat{1}]$ admits a recursive atom ordering in which the atoms of $[a_j, \hat{1}]$ that come first in the ordering are those that cover some a_i where $i < j$.
2. For all $i < j$, if $a_i, a_j < y$ then there is a $k < j$ and an element z such that z covers a_k and a_j, and $z \leq y$.

In [6] Björner and Wachs establish the following result:

Lemma 5.3. (Proposition 2.3 & Theorem 3.2 in [6]) *Any bounded, graded poset that admits a recursive atom ordering is shellable, and therefore Cohen-Macaulay.*

Recall that a poset P is *bounded* if there is a minimal element $\hat{0}$ that is less than all other elements in P, as well as a maximal element $\hat{1}$ that is greater than all other elements in P.

A particularly strong result is available for totally semimodular posets.

Definition 5.4 (Semimodular)**.** A graded poset P is *semimodular* if it is bounded and whenever two distinct elements $u, v \in P$ cover $x \in P$ there is a $z \in P$ which covers both u and v. A graded poset P is *totally semimodular* if it is bounded and all intervals $[x, y]$ are semimodular.

Lemma 5.5 (Theorem 5.1 in [6])**.** *A graded poset P is totally semimodular if and only if for every interval $[x, y]$ of P, every ordering of the atoms in $[x, y]$ is a recurvsive atom ordering.*

The key technical result in this section is that the poset Z_n, defined below, has a recursive atom ordering.

Definition 5.6 (Z_n)**.** Let Z_n be the poset obtained from W_n by taking the dual poset and adding a minimal element $\hat{0}$. (So $\hat{0} < T$ for all $T \in Z_n \backslash \{\hat{0}\}$.) Since the dual of W_n already has a unique maximal element, namely the nuclear vertex, the poset Z_n is a bounded, graded poset. We adopt the standard notation and denote the nuclear vertex, thought of as an element of Z_n, by $\hat{1}$. Thus, in Z_n one has $T < T'$ if and only if $T = \hat{0}$, or $T \neq \hat{0}$ and one can convert T to T' by a sequence of folds. The poset Z_4 is shown in Figure 4.

Lemma 5.7. *For each $T \in Z_n$ which is not equal to $\hat{0}$, the interval $[T, \hat{1}]$ is totally semimodular.*

Proof. Let T' be any tree in $[T, \hat{1}]$ and let S and S' be two trees in the interval that cover T'. If S and S' resulted from folds in T' that occur at different vertices of

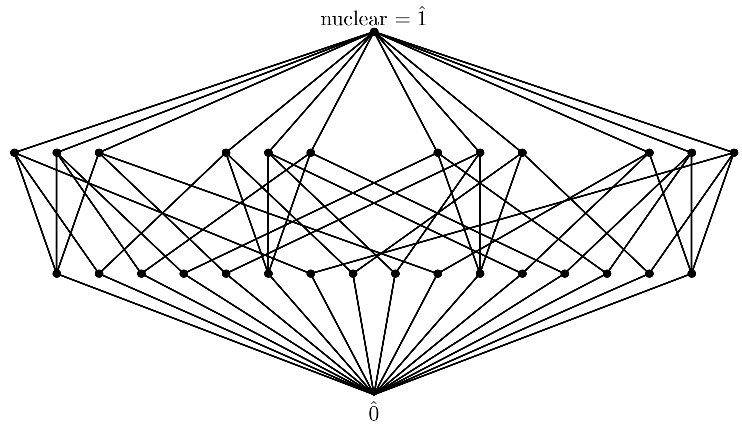

FIGURE 4. The poset Z_4.

T', then it is clear that there is a tree $Z \in [T, \hat{1}]$ covering both S and S' since the set of edges incident to a vertex v_i are identical before and after a fold at a vertex v_j, so long as $i \neq j$. If S and S' result from folds at the same vertex, then either two distinct pairs of edges are being folded together, or the folds have an edge in common. In either event, there is a tree Z covering S and S' gotten by performing both folds. □

It follows from Lemma 5.7 and Lemma 5.5 that any ordering of the atoms in Z_n gives a recursive atom ordering on intervals. Thus the first condition can always be satisfied, and we can turn our attention to the second condition. The recursive atom ordering of the atoms of Z_n will be based on the following partial ordering of $[n]$-trees.

Definition 5.8 (Depth ordering)**.** Every $[n]$-tree can be rooted at 1 and drawn in the standard fashion for rooted trees as in Figure 5. We say the labelled vertices of T immediately below 1 in such a drawing are at "level 1". If we think of an $[n]$-tree T as a metric tree where each edge has length 1/2, then the vertices at level 1 are precisely those that are a distance 1 from the vertex v_1, i.e. the vertex labelled 1. In general, the *level* of a labelled vertex v is simply $d(v_1, v)$.

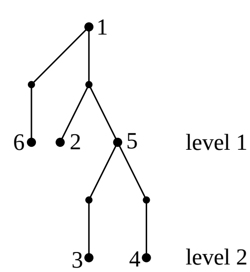

FIGURE 5. Depth of vertices.

The *depth* of an $[n]$-tree T is its maximum level, or equivalently, the radius of T with center v_1. Our partial ordering is given first by depth "d", then by the number of labelled vertices at level d, then by the number of vertices at level d-1, and so on. Finally, we extend this partial ordering of $[n]$-trees to a total ordering arbitrarily. We denote this total ordering by $<_{\text{depth}}$.

Definition 5.9 (Drops and splits)**.** Let T be any $[n]$-tree. We divide the unfoldings of T into two groups: "drops" and "splits". The terminology is transparent when one views T as being rooted at the vertex labelled 1. Since a fold identifies two edges at a folding vertex v, an unfold divides an edge incident to the unfolding vertex v. A *split* is any unfolding where the edge being divided hangs below the unfolding vertex. A *drop* is any unfolding where the edge being divided is above the unfolding vertex. A drop based at the unfolding vertex v will move certain labelled vertices at the same level as v down one level. The foldings which reverse a drop or a split will be called *lifts* and *merges* respectively. See Figure 6 and Figure 7 for illustrations.

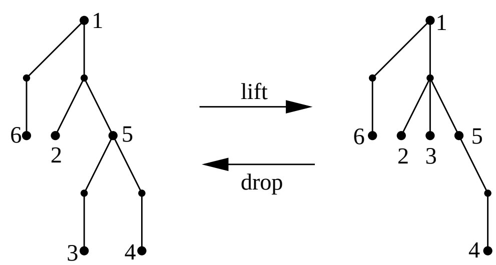

FIGURE 6. Dropping and lifting.

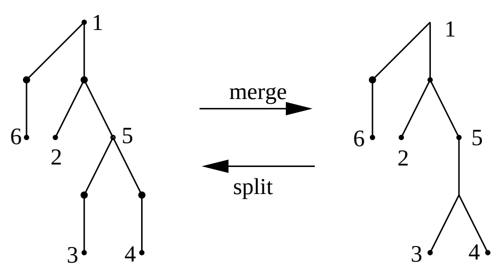

FIGURE 7. Splitting and merging.

The following lemma is immediate from the definitions.

Lemma 5.10. *If S results from T by a drop, then T lies strictly below S in the depth partial order.*

The atoms of Z_n consist of all $[n]$-trees where the unlabelled vertices have valence 2. Equivalently, they are the trees of rank $(n-2)$. Notice that the minimal atom in the total order $<_{\text{depth}}$ is uniquely determined by the depth partial order; it is the tree with central vertex labelled 1, all unlabeled vertices with valence 2 and all other labelled vertices at level 1. It is the unique atom which results by repeatedly splitting the nuclear vertex.

Corollary 5.11. *Given any $[n]$-tree T there is a unique atom $m(T)$ such that T can be unfolded to $m(T)$ and $m(T)$ and T are incomparable in the depth partial order. Further, if S is any other atom such that T can be unfolded to S, then $m(T) <_{depth} S$.*

Proof. You can get from any $[n]$-tree to an atom using only splits and splits do not change the data used to define the depth partial order. Further, if one does all possible splits in any order, then they will always end up at the same tree, which will be denoted $m(T)$. On the other hand, if the sequence of unfoldings from T to an atom S contains a drop, then T is less than S in the depth partial order by Lemma 5.10, and therefore $m(T) <_{\text{depth}} S$. □

Theorem 5.12. *Any total ordering of the atoms of Z_n that is compatible with the depth ordering gives a recursive atom ordering of Z_n.*

Proof. Let T_i and T_j be two atoms of Z_n with T_i less than T_j in the total order $<_{\text{depth}}$ and suppose both $[n]$-trees can be folded up to an $[n]$-tree T. To establish condition 2 in the definition of a recursive atom ordering, we need to find an atom T_k such that $T_k <_{\text{depth}} T_j$ and an element $Z \leq T$ that is a cover of both T_k and T_j.

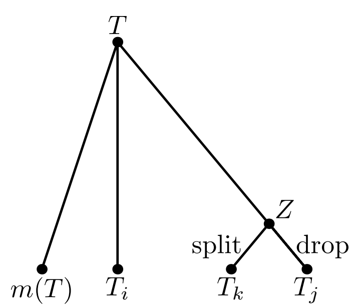

FIGURE 8. Illustration of the proof.

Since T_i is strictly below T_j in the total order $<_{\text{depth}}$, and by Corollary 5.11 either $m(T) = T_i$ or $m(T) <_{\text{depth}} T_i$, we know that $m(T) <_{\text{depth}} T_j$. This shows that the sequence of unfoldings from T down to T_j must contain at least one drop. Reverse this sequence of unfoldings and consider the sequence of foldings from T_j up to T. If a merge is followed by a lift, then one can replace this pair of folds by a pair beginning with a lift. The argument is essentially that of Lemma 5.7. The only interesting case is when the folding vertex v is the same for both the merge and lift, and one first merges two edges e' and e'' into one edge e and then lifts e. In this case the same result occurs if one would first lift e' and then lift e''. As a result, we can choose the sequence of folds so that the first fold at T_j is a lift. Let Z be the result of this lift and let T_k be the atom $m(Z)$. By Lemma 5.10 $T_k <_{\text{depth}} T_j$, Z is a common cover of T_j and T_k, and by construction $Z \leq T$ in Z_n.

This shows that the total ordering on the atoms satisfies condition 2, and as was mentioned earlier, by Lemma 5.7 and Lemma 5.5 it can be recursively extended upwards to satisfy condition 1. This completes the proof. □

Theorem 5.13. *The Whitehead posets W_n are Cohen-Macaulay.*

Proof. By Theorem 5.12, Z_n has a recursive atom ordering and thus is Cohen-Macaulay by Lemma 5.3. The maximal chains of Z_n are the same as the maximal chains of $\widehat{W}_n$ where $\widehat{W}_n$ is simply the Whitehead poset W_n with a single maximal element $\hat{1}$ attached. Thus any shelling of Z_n induces a shelling of $\widehat{W}_n$ hence $\widehat{W}_n$ is Cohen-Macaulay. The simplicial complex $|W_n|$ is the link of the vertex $\hat{1} \in |\widehat{W}_n|$. The theorem follows since the link of any vertex in a Cohen-Macaulay complex is Cohen-Macaulay. □

6. Further questions

In addition to being realizable as a natural subgroup of $\mathrm{Aut}(F_n)$, the group $\mathrm{P}\Sigma_n$ arises as a motion group. The pure braid group can be thought of as the group of motions of n points in the plane; $\mathrm{P}\Sigma_n$ consists of the motions of the trivial n component link in S^3. This is discussed in [10], where it is pointed out that the group of motions of n unknotted, unlinked k-spheres in S^{k+2} can also be represented in $\mathrm{Aut}(F_n)$. We can denote these groups as $\mathrm{P}\Sigma_n^k$, and ask:

Question 6.1. Is there a natural description of the images of the $\mathrm{P}\Sigma_n^k$ in $\mathrm{Aut}(F_n)$? Is the representation even faithful? What is the dimension of these groups? Are they all duality groups?

Perhaps an even more elementary question would be

Question 6.2. Are the groups of motions of non-trivial links in S^3 virtual duality groups, for all non-trivial links? In other words, was the assumption that we were working with the trivial n component link necessary?

The analogy between $\mathrm{P}\Sigma_n$ and the pure braid group has led a number of people to wonder whether duality could be established for these groups as in the pure braid case. For the pure braid groups there is a natural map $PB_n \twoheadrightarrow PB_{n-1}$ which is obtained by ignoring the n^{th} strand. Since the kernel of this map is free, an induction shows that PB_n is poly-free. For example, PB_3 is F_2-by-$\mathbb{Z}$, and PB_4 is F_3-by-(F_2-by-$\mathbb{Z}$), etc.. Thus since a duality-by-duality group is a duality group, of dimension equal to the sum of the dimensions of the kernel and quotient, PB_n is an $(n-1)$-dimensional duality group.

Regretably, there are significant obstructions to extending this argument to $\mathrm{P}\Sigma_n$. The group $\mathrm{P}\Sigma_n$ is generated by automorphisms α_{ij} where $\alpha_{ij}(x_i) = x_j^{-1}x_ix_j$ and $\alpha_{ij}(x_k) = x_k$, if $(k \neq i)$. (Such automorphisms correspond to moving one loop through another.) The kernel of the natural map $\mathrm{P}\Sigma_n \twoheadrightarrow \mathrm{P}\Sigma_{n-1}$ is then the normal subgroup generated by the automorphisms α_{in} and α_{ni}, where $i \in [n-1]$. In particular, this subgroup is not free, since $\langle\{\alpha_{in} \mid i \in [n-1]\}\rangle$ is free abelian of rank $n-1$. Further, $\langle\{\alpha_{in}, \alpha_{ni} \mid i \in [n-1]\}\rangle$ is not a normal subgroup of $\mathrm{P}\Sigma_n$. Thus it is not immediately apparent that the kernel of $\mathrm{P}\Sigma_n \twoheadrightarrow \mathrm{P}\Sigma_{n-1}$ is even finitely generated.

Question 6.3. What can be said about the kernel of the map $\mathrm{P}\Sigma_n \twoheadrightarrow \mathrm{P}\Sigma_{n-1}$?

Finally, $\mathrm{P}\Sigma_n$ has a presentation in which all the defining relations are commutators (see for example [13]). This makes it appear likely that $\mathrm{P}\Sigma_n$ might be a CAT(0) group. (This is similar to the belief that braid groups might be CAT(0) groups.) However, $\mathrm{P}\Sigma_3$ is a 2-dimensional group, and careful computations show that its presentation 2-complex does not support a CAT(0) metric.

Question 6.4. Is $P\Sigma_n$ a CAT(0) group?

Finally, we remind the reader that Gutiérrez and Krstić have established that $P\Sigma_n$ has a regular language of normal forms, and they have asked if moreover $P\Sigma_n$ is (bi)automatic [11].

References

[1] M. Bestvina and M. Feighn, The topology at infinity of $\mathrm{Out}(F_n)$, *Invent. Math.* **140** (2000) 651–692.

[2] R. Bieri, Homological Dimension of Discrete Groups, 2nd ed., Queen Mary College Mathematics Notes, 1981.

[3] R. Bieri and B. Eckmann, Groups with homological duality generalizing Poincaré duality, *Invent. Math.* **20** (1973) 103–124.

[4] B. Bogley and S. Krstić, In progress.

[5] A. Björner, Shellable and Cohen-Macaulay partially ordered sets, *Trans. Amer. Math. Soc.* **260** (1980), 159–183.

[6] A. Björner and M. Wachs, On lexicographically shellable posets, *TAMS*, **277** (1983) 323–341

[7] K.S. Brown, *Cohomology of Groups*, Springer-Verlag, 1982.

[8] K.S. Brown and J. Meier, Improper actions and higher connectivity at infinity, *Comment. Math. Helv.* **75** (2000) 171–188.

[9] D.J. Collins, Cohomological dimension and symmetric automorphisms of a free group, *Comment. Math. Helv.* **64** (1989) 44–61.

[10] D.L. Goldsmith, The theory of motion groups, *Mich. Math. J.* **28** (1981) 3–17.

[11] M. Gutiérrez and S. Krstić, Normal forms for basis-conjugating automorphisms of a free group, *IJAC* **8** (1998) 631–669.

[12] Darryl McCullough and Andy Miller, Symmetric Automorphisms of Free Products, *Memoirs of the AMS* #582, 1996.

[13] L. Orlandi-Korner, The Bieri-Neumann-Strebel invariant for basis-conjugating automorphisms of free groups, *Proc. Amer. Math. Soc.* **128** (2000) 1257–1262.

Department of Mathematics, University of Oklahoma, Norman, OK 73019
E-mail address: `nbrady@math.ou.edu`

Department of Mathematics, Texas A & M University, College Station, TX 77843
E-mail address: `jon.mccammond@math.tamu.edu`

Department of Mathematics, Lafayette College, Easton, PA 18042
E-mail address: `meierj@lafayette.edu`

Department of Mathematics, University of Oklahoma, Norman, OK 73019
E-mail address: `amiller@ou.edu`