Canad. J. Math. 2026, pp. 1–27
http://dx.doi.org/10.4153/S0008414X26102089
© The Author(s), 2026. Published by Cambridge University Press on behalf of Canadian Mathematical Society. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Dense and empty BNSR-invariants of the McCool groups

Mikhail Ershov and Matthew C. B. Zaremsky

Abstract. An automorphism of the free group F_n is called *pure symmetric* if it sends each generator to a conjugate of itself. The group PSAut_n of all pure symmetric automorphisms and its quotient PSOut_n by the group of inner automorphisms are called the *McCool groups*. In this article, we prove that every BNSR-invariant Σ^m of a McCool group is either dense or empty in the character sphere, and we characterize precisely when each situation occurs. Our techniques involve understanding higher generation properties of abelian subgroups of McCool groups, coming from the McCullough–Miller space. We also investigate further properties of the second invariant Σ^2 for McCool groups using a general criterion due to Meinert for a character to lie in Σ^2.

1 Introduction

Let G be a finitely generated group. The *BNS-invariant* $\Sigma(G)$ is a fundamental geometric object introduced in the seminal 1987 paper of Bieri, Neumann, and Strebel [BNS87]. In 1988, two families of its higher-dimensional generalizations were defined: *homological BNSR-invariants* $\Sigma^m(G,R)$, $m\in\mathbb{N}$, where R is a commutative unital ring, with $\Sigma^1(G,\mathbb{Z})=\Sigma(G)$, studied by Bieri and Renz in [BR88], and *homotopical BNSR-invariants* $\Sigma^m(G)$, $m\in\mathbb{N}$, with $\Sigma^1(G)=\Sigma(G)$, studied by Renz in [Ren88]. In this article, we will only deal with homotopical invariants and refer to them simply as *BNSR-invariants*. The mth invariant $\Sigma^m(G)$ is defined whenever G is of *type* F_m, meaning G admits a classifying space with finite m-skeleton. In general, the BNSR-invariants of G form a descending chain $\Sigma^1(G)\supseteq\Sigma^2(G)\supseteq\cdots$ of open subsets of the sphere $\mathbb{S}(G)=S^d$, where $d=\dim\mathrm{Hom}(G,\mathbb{R})-1$. Among other things, $\Sigma^m(G)$ completely determines which coabelian subgroups of G are of type F_m (we call a subgroup of G *coabelian* if it contains the commutator subgroup $[G,G]$); in particular, $[G,G]$ itself is of type F_m if and only if $\Sigma^m(G)=\mathbb{S}(G)$.

While the BNS-invariant $\Sigma(G)$ has been fully computed in a number of interesting cases, the class of groups for which all BNSR-invariants $\Sigma^m(G)$ are defined and completely determined for all m is much smaller. One prominent class for which the latter has been accomplished is right-angled Artin groups (RAAGs) [MMV98], and it is natural to ask what can be said about BNSR-invariants for groups which are sufficiently

Received by the editors June 18, 2025; revised January 22, 2026; accepted February 2, 2026.
Published online on Cambridge Core February 13, 2026.
M.Z. was supported for part of the time spent on this project by Simons Grant No. 635763.
AMS subject classification: **20F65**, 57M07.
Keywords: BNSR-invariant, automorphisms of free groups, McCool group, loop braid group, finiteness properties.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

close to RAAGs in some sense. Examples of such groups where some progress has been made include Artin groups in general [BGCAMP22, EFMP24, MMV01] and Bestvina–Brady groups [KM22]. In this article, we will focus on the computation of BNSR-invariants for yet another collection of groups similar to RAAGs, which are frequently called *McCool groups*.

Let F_n be the free group of rank n, with a fixed basis $\{x_1, \dots, x_n\}$. An automorphism of F_n is called *pure symmetric* if it sends each x_i to a conjugate of itself. The subgroup PSAut_n of the automorphism group $\mathrm{Aut}(F_n)$ consisting of all pure symmetric automorphisms arises in many contexts and has several different names. It is sometimes simply called the *pure symmetric automorphism group*, sometimes the *McCool group*, in reference to seminal work of McCool [McC86], and sometimes the *pure loop braid group*, for its appearance as the group of motions of n unlinked distinguishable circles in 3-space (see, e.g., [BH13, BWC07] for some important applications of these groups and [Dam17] for an overview of all their different guises). The "outer" version PSOut_n of PSAut_n is the quotient of PSAut_n by the group of inner automorphisms, and we will collectively refer to all the PSAut_n and PSOut_n as *McCool groups*.

At a very superficial level, McCool groups are similar to RAAGs because they admit simple presentations by generators and relations where all relations are of the form $[u, v] = 1$ for some words u and v, and in the majority of relations, u and v are both generators. However, there are much deeper analogies between these groups, which make many techniques and ideas from [MMV98] applicable to McCool groups.

The first BNS-invariant for the McCool groups has been completely determined by Orlandi-Korner in [OK00]. In [Zar18], the second author computed some pieces of the higher BNSR-invariants $\Sigma^m(\mathrm{PSAut}_n)$ and conjectured that each $\Sigma^m(\mathrm{PSAut}_n)$ is either dense (in the respective sphere) or empty, depending on whether $m \leq n-2$ or not [Zar18, Remark 4.24]. The main theorem of this article confirms this conjecture and also establishes the analogous result for the groups PSOut_n.

Theorem 1.1 *For all $n \geq 2$, the following hold:*

(a) $\Sigma^{n-2}(\mathrm{PSAut}_n)$ *is dense in* $\mathbb{S}(\mathrm{PSAut}_n)$.
(b) $\Sigma^{n-1}(\mathrm{PSAut}_n) = \varnothing$.

For all $n \geq 3$, the following hold:

(c) $\Sigma^{n-3}(\mathrm{PSOut}_n)$ *is dense in* $\mathbb{S}(\mathrm{PSOut}_n)$.
(d) $\Sigma^{n-2}(\mathrm{PSOut}_n) = \varnothing$.

Hence, $\Sigma^k(\mathrm{PSAut}_n)$ is dense for all $0 \leq k \leq n-2$ and empty for all $k \geq n-1$, and $\Sigma^k(\mathrm{PSOut}_n)$ is dense for all $0 \leq k \leq n-3$ and empty for all $k \geq n-2$.

Remark 1.2

(i) Since we always have $\mathbb{S}(G) = \Sigma^0(G) \supseteq \Sigma^1(G) \supseteq \cdots$, the "hence" part is immediate from parts (a)–(d).
(ii) All assertions of Theorem 1.1 were previously known for $n \leq 3$. Indeed, $\Sigma^1(\mathrm{PSAut}_2)$ is empty and $\Sigma^1(\mathrm{PSAut}_3)$ is dense in $\mathbb{S}(\mathrm{PSAut}_3)$ by the main theorem of [OK00], and the proof of that theorem also easily implies that $\Sigma^1(\mathrm{PSOut}_3)$ is empty. Finally, the equality $\Sigma^2(\mathrm{PSAut}_3) = \varnothing$ was established by the second author in [Zar18].

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

While Theorem 1.1 naturally extends the results and confirms the aforementioned conjecture from [Zar18], the techniques used in the proofs are very different, as we will see later.

A consequence of Theorem 1.1, together with Theorem 2.3 below, is that if $\chi\colon \mathrm{PSAut}_n \to \mathbb{Z}$ is a "generic" discrete character of PSAut_n, then the kernel of χ is of type F_{n-2} but not F_{n-1}. Similarly, the kernel of a "generic" discrete character of PSOut_n is of type F_{n-3} but not F_{n-2}. We should also mention that our proof of parts (b) and (d) actually shows the stronger statement that $\Sigma^{n-1}(\mathrm{PSAut}_n;\mathbb{Z}) = \varnothing$ and $\Sigma^{n-2}(\mathrm{PSOut}_n;\mathbb{Z}) = \varnothing$, where these are the homological BNSR-invariants.

In order to prove parts (b) and (d), about certain BNSR-invariants being empty, we prove a rather easy (and somewhat "folklore"), but very general, result about groups of type F (meaning groups admitting a finite classifying space) with nonzero Euler characteristic – see Proposition 3.1 which asserts that if G is a group of type F whose BNSR-invariants are symmetric (Definition 2.4) and whose Euler characteristic is nonzero, then $\Sigma^{\infty}(G) = \varnothing$. Since for a group of type F, we have $\Sigma^{\infty}(G) = \Sigma^{m}(G)$, where m is the dimension of a finite classifying space for G, parts (b) and (d) follow easily from known facts about the McCool groups.

Proving parts (a) and (c), about certain BNSR-invariants being dense, requires quite a lot more effort. The key is a criterion due to Meier, Meinert, and Van Wyk [MMV98] that relates the BNSR-invariants of a group to the BNSR-invariants of its subgroups, using the notion of higher generation of groups by families of subgroups introduced by Abels and Holz in [AH93] (see Theorem 4.3). The Meier–Meinert–Van Wyk criterion involves a number of hypotheses on the group, and in order to apply it to the McCool groups, we make use of the Whitehead posets and McCullough–Miller spaces (see Definitions 4.6 and 4.16) and the fact that the groups are highly generated by abelian subgroups (see Proposition 4.17). In a recent paper [AGMP25], a generalization of McCullough–Miller space for RAAGs is constructed with all these desired properties, so it would be interesting to use this to analyze the BNSR-invariants of the analogs of the McCool groups for RAAGs.

Given that PSAut_n is sometimes called the "pure loop braid group," it is natural to ask if one could prove similar results about the pure braid groups PB_n. Since PB_n splits as $(\mathrm{PB}_n/Z(\mathrm{PB}_n)) \times \mathbb{Z}$ [FM12, Section 9.3], $\Sigma^{\infty}(\mathrm{PB}_n)$ is dense in its character sphere for trivial reasons, but an interesting problem is understanding when $\Sigma^{m}(\mathrm{PB}_n/Z(\mathrm{PB}_n))$ is dense or empty. In fact, $\mathrm{PB}_n/Z(\mathrm{PB}_n)$ is isomorphic to the pure mapping class group of the $(n+1)$-punctured sphere (equivalently the n-punctured plane) [FM12, Section 9.3], and it is already known that $\Sigma^{n-2}(\mathrm{PB}_n/Z(\mathrm{PB}_n)) = \varnothing$ for all $n \geq 3$ [HK24, Example 7.3] (when $n < 3$, the group is trivial). Thus, a natural conjecture is that $\Sigma^{n-3}(\mathrm{PB}_n/Z(\mathrm{PB}_n))$ is dense in its character sphere. This is trivially true for $n = 3$ and can be confirmed for $n = 4$ thanks to the computation of $\Sigma^{1}(\mathrm{PB}_4)$ in [KMM15]. An impediment to proving this conjecture for larger n is a lack of a known analog for PB_n of the McCullough–Miller space for PSAut_n and PSOut_n and, more generally, a lack of a known highly generating family of abelian subgroups. If one could find a family of abelian subgroups of PB_n that is ∞-generating in the sense of Definition 4.1, then proving that $\Sigma^{n-3}(\mathrm{PB}_n/Z(\mathrm{PB}_n))$ is dense would probably be within reach. We should mention that some results about highly generating families

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

for PB_n are in the Appendix of [BFM+16], but the degree of high generation is not large enough to tell us anything about $\Sigma^{n-3}(\mathrm{PB}_n/Z(\mathrm{PB}_n))$.

We also prove further results about the second BNSR-invariant Σ^2 for McCool groups. In dimension 2, an underutilized result from the Ph.D. thesis of Meinert [Mei93] leads to an easy-to-check sufficient condition for a character class to lie in $\Sigma^2(Q)$, where Q is a quotient of a well-understood group G (see Theorem 5.4 for the general criterion and Proposition 5.3 for the specific application to McCool groups). Since Meinert's thesis is not easily accessible, we also provide a proof of Meinert's theorem (see Theorem A.1 in Appendix A).

This article is organized as follows. In Section 2, we recall some background material on BNSR-invariants and McCool groups and set the stage for inspecting BNSR-invariants of McCool groups. In Section 3, we prove parts (b) and (d) of Theorem 1.1, about BNSR-invariants being empty. In Section 4, we prove parts (a) and (c), about BNSR-invariants being dense. In Section 5, we prove our further results about Σ^2, stemming from Meinert's criterion, and then in Appendix A, we provide a proof of Meinert's criterion.

2 Preliminaries

In this section, we recall some background material on BNSR-invariants in general and on the groups PSAut_n and PSOut_n that will be our main objects of interest.

2.1 BNSR-invariants

In this section, we recall the definition and some important properties of the Bieri–Neumann–Strebel–Renz (BNSR) invariants of a group. The invariants are somewhat difficult to define and usually quite difficult to compute, but reveal a wealth of information about the group. Let us gradually build up their definition.

Recall that a *classifying space* for a group G is a connected CW-complex whose fundamental group is isomorphic to G and whose homotopy groups vanish in all dimensions higher than 1. Equivalently, it is a quotient of a connected CW-complex whose homotopy groups in positive dimensions are trivial by a free cellular action of G.

A group is said to be of *type* F_n if it admits a classifying space with a finite n-skeleton. Every group is of type F_0, finite generation is equivalent to type F_1, and finite presentability is equivalent to type F_2. A group is of *type* F_∞ if it is of type F_n for all n, and if a group has a finite classifying space, it is said to be of *type F*.

Throughout the article, by a *character* of a group G we will mean a homomorphism $\chi\colon G \to \mathbb{R}$ from G to the additive group of real numbers. Two characters are *equivalent* if they differ by multiplication by a positive scalar. The equivalence classes $[\chi]$ of nontrivial characters of G form the *character sphere* of G, denoted $\mathbb{S}(G)$.

This sphere should be viewed as the boundary of the Euclidean vector space $\mathrm{Hom}(G,\mathbb{R})$. If this space is d-dimensional, then the character sphere of G is $(d-1)$-dimensional. BNSR-invariants of G are certain subsets of $\mathbb{S}(G)$, which can be defined in several slightly different ways. The definition below is close, but not identical, to the original definition from [Ren88].

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

Let G be a group of type F_n, so by definition it has a classifying space with finite n-skeleton. Choose any such classifying space K, let $\widetilde{K}$ be the universal cover of K, and let X be the n-skeleton of $\widetilde{K}$, with the canonical G-action. It is not hard to show that for any nontrivial character χ of G there exists a map $h_\chi\colon X \to \mathbb{R}$ (called a *height function*) such that $h_\chi(g.x) = \chi(g) + h_\chi(x)$ for all $g \in G$ and $x \in X$. For $t \in \mathbb{R}$, define $X^{h_\chi \geq t}$ to be the subcomplex of X consisting of all cells whose vertices lie in $h_\chi^{-1}([t, \infty))$.

Definition 2.1 (BNSR-invariants) Let n, G, and X be as above. For any $0 \leq m \leq n$, the mth *BNSR invariant* $\Sigma^m(G)$ is the subset of $\mathbb{S}(G)$ consisting of all $[\chi]$ for which the filtration $X^{h_\chi \geq t}$ of X is *essentially* $(m-1)$*-connected*, meaning that for every t, there exists $u \leq t$ such that the inclusion $X^{h_\chi \geq t} \to X^{h_\chi \geq u}$ induces the trivial map in all homotopy groups up to dimension $m-1$.

Note that by construction, X in Definition 2.1 is $(n-1)$-connected, so all k-spheres in X can be filled for any $k \leq n-1$, but what the condition $[\chi] \in \Sigma^m(G)$ measures is whether they can be filled without causing the χ value to drop too far in the negative direction somewhere along the filling disk. This definition is independent of the choices of X and h_χ (see, e.g., [BG03, Theorem 12.1]).

Remark 2.2

(a) One can compute BNSR-invariants using more general G-complexes. For instance, one can take X to be any $(n-1)$-connected CW-complex on which G acts properly, cocompactly, and cellularly (see, e.g., [Zar18, Definition 1.1]); for an even more general condition, see [Bux04, Definition 8.1].

(b) For any X and χ as in Definition 2.1, to prove that the filtration $X^{h_\chi \geq t}$ is essentially $(m-1)$-connected it suffices to check the corresponding condition just for $t = 0$. This follows easily from the fact that $g.X^{h_\chi \geq t} = X^{h_\chi \geq t + \chi(g)}$ for all $g \in G$ and $t \in \mathbb{R}$.

(c) The original definition from [Ren88] was as follows: $[\chi] \in \Sigma^m(G)$ if there exists some X as in Definition 2.1 such that the subcomplex $X^{h_\chi \geq t}$ is $(m-1)$-connected for all $t \in \mathbb{R}$. Note that unlike essential $(m-1)$-connectivity, this condition does depend on the choice on X. The equivalence of this definition and Definition 2.1 was established already in [Ren88] (see also [Mei97, Section 2] and [LI24, Section 2]).

If $n = 1$, a natural choice for X in Definition 2.1 is the Cayley graph $\mathrm{Cay}(G, S)$ with respect to some finite generating set S. Likewise, if $n = 2$, so that G is finitely presented, we can take X to be the Cayley complex associated with a finite presentation of G.

If $n = 1$ and $X = \mathrm{Cay}(G, S)$ for a finite generating set S, there is no difference between connectivity and essential connectivity (for any S); in fact in this case, the following are equivalent for a nonzero character χ:

(a) $[\chi] \in \Sigma^1(G)$;
(b) $X^{h_\chi \geq t}$ is connected for all t;
(c) $X^{h_\chi \geq t}$ is connected for some t.

It is clear that (b) implies (a). That (c) implies (b) is proved, for example, in [Str12, Lemma A2.9]. Finally, to prove that (a) implies (c) note that for any $u \leq t$, any vertex of $X^{h_\chi \geq u}$ is connected by a path in $X^{h_\chi \geq u}$ to a vertex in $X^{h_\chi \geq t}$ (just choose $s \in S^{\pm 1}$ with $\chi(s) > 0$ and multiply by s^n on the right for sufficiently large n); hence, if $X^{h_\chi \geq t}$ is contained in a connected subset of $X^{h_\chi \geq u}$, then the entire $X^{h_\chi \geq u}$ is connected.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

If G is of type F_∞, then $\Sigma^m(G)$ is defined for all m, and we have a nested sequence

$$\mathbb{S}(G) \supseteq \Sigma^1(G) \supseteq \Sigma^2(G) \supseteq \cdots.$$

The intersection of all the $\Sigma^m(G)$ is denoted $\Sigma^\infty(G)$. Note that if G is of type F with a finite classifying space of dimension n, then $\Sigma^n(G) = \Sigma^\infty(G)$ (e.g., this follows from [Zar17, Citation 1.5]). Thus, for groups of type F, the above sequence stabilizes.

An important property of BNSR-invariants, proved in [BNS87, Ren88], is that $\Sigma^m(G)$ is open in $\mathbb{S}(G)$ for all $m < \infty$. Another key property is that $\Sigma^m(G)$ is invariant under the action of $\mathrm{Aut}(G)$, that is, if α is an automorphism of G, then for a character χ, we have $[\chi \circ \alpha] \in \Sigma^m(G)$ if and only if $[\chi] \in \Sigma^m(G)$. This easily follows from the fact that $\Sigma^m(G)$ in Definition 2.1 does not depend on the choice of G-complex X or the choice of h_χ, so in particular, one can keep the same complex X, but "twist" the action by α.

The main application of BNSR-invariants is the following theorem of Renz [Ren88, Satz C] (see also Theorem 1.1 of [BGK10]).

Theorem 2.3 [Ren88, Satz C] *Let G be a group of type F_n and H a subgroup of G containing the commutator subgroup $[G,G]$. Then H is of type F_n if and only if $[\chi] \in \Sigma^n(G)$ for all $0 \neq \chi \in \mathrm{Hom}(G,\mathbb{R})$ such that $\chi(H) = 0$.*

For example, $[G,G]$ itself is of type F_n if and only if $\Sigma^n(G) = \mathbb{S}(G)$. At the other extreme in a sense, if $0 \neq \chi\colon G \to \mathbb{R}$ is *discrete*, meaning its image is cyclic, then $\ker(\chi)$ is of type F_n if and only if $[\chi], [-\chi] \in \Sigma^n(G)$. In practice, many groups have BNSR-invariants where one does not need to worry about checking both χ and $-\chi$, since they cannot have different behavior. Let us write

$$-\Sigma^m(G) = \{[-\chi] \mid [\chi] \in \Sigma^m(G)\},$$

and make the following definition.

Definition 2.4 (Symmetric) We say that the BNSR-invariants of a group G are *symmetric* if $\Sigma^m(G) = -\Sigma^m(G)$ for all m. For example, if there is an automorphism ω of G such that $\chi \circ \omega = -\chi$ for all χ, then the BNSR-invariants of G are symmetric.

2.2 McCool groups

In this section, we introduce our groups of interest. Let F_n be the free group of rank n, with a fixed basis $\{x_1, \ldots, x_n\}$. Consider the group $\mathrm{Aut}(F_n)$ of automorphisms of F_n. Denote by $\mathrm{Inn}(F_n)$ the subgroup of inner automorphisms, and let $\mathrm{Out}(F_n) = \mathrm{Aut}(F_n)/\mathrm{Inn}(F_n)$ be the group of outer automorphisms. An automorphism of F_n is called *symmetric* if it sends each x_i to a conjugate of some x_j, and *pure symmetric* if it sends each x_i to a conjugate of x_i. Denote the group of all symmetric automorphisms of F_n by SAut_n, and the group of all pure symmetric automorphisms of F_n by PSAut_n. Since all inner automorphisms are pure symmetric, we have

$$\mathrm{Inn}(F_n) \leq \mathrm{PSAut}_n \leq \mathrm{SAut}_n \leq \mathrm{Aut}(F_n).$$

Taking the quotient by $\mathrm{Inn}(F_n)$, we obtain "outer" versions:

$$\mathrm{PSOut}_n \leq \mathrm{SOut}_n \leq \mathrm{Out}(F_n).$$

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

In [McC86], McCool exhibited a natural finite presentation for PSAut_n. There is a generator α_{ij} for each $1 \le i, j \le n$ with $i \ne j$, and the defining relations are as follows:

- $[\alpha_{ij}, \alpha_{k\ell}] = 1$ for i, j, k, ℓ distinct;
- $[\alpha_{ij}, \alpha_{kj}] = 1$ for i, j, k distinct;
- $[\alpha_{ij}\alpha_{kj}, \alpha_{ik}] = 1$ for i, j, k distinct.

The generator α_{ij} is the automorphism of F_n sending x_i to $x_i^{x_j} = x_j^{-1}x_i x_j$ and fixing all x_k for $k \ne i$. We will generally refer to PSAut_n and PSOut_n as *McCool groups*. Note that PSAut_1 is trivial, so when talking about McCool groups we will often implicitly assume $n \ge 2$, which will avoid some pathologies.

Thanks to the second relation, if $I \subseteq \{1, \dots, n\}$ and $j \in \{1, \dots, n\} \backslash I$, then α_{ij} and $\alpha_{i'j}$ commute for all $i, i' \in I$. In particular, it makes sense to define

$$\alpha_{I,j} = \prod_{i \in I} \alpha_{ij},$$

that is, the order of the product does not matter. The automorphisms $\alpha_{I,j}$ are called the *Whitehead generators* of PSAut_n. We say that $\alpha_{I,j}$ is *based* at j. Note that $\mathrm{Inn}(F_n)$ is the subgroup of PSAut_n generated by all the $\alpha_{[n]\backslash\{j\},j}$ for $1 \le j \le n$, where

$$[n] = \{1, \dots, n\}.$$

We will use the following basic relations between the Whitehead generators, which are all clear from the definition.

Lemma 2.5 *The following hold:*

(i) $\alpha_{I,j}\,\alpha_{K,j} = \alpha_{I\cup K,j}$ *whenever* I *and* K *are disjoint;*
(ii) $[\alpha_{I,j}, \alpha_{K,j}] = 1$;
(iii) $[\alpha_{I,j}, \alpha_{K,m}] = 1$ *if either* $I \cup \{j\}$ *and* $K \cup \{m\}$ *are disjoint or* $I \cup \{j\} \subseteq K$.

The McCool groups have some important topological properties that we will need to use later.

Proposition 2.6 *The group* PSAut_n *is torsion-free and admits a free simplicial cocompact action on an* $(n-1)$*-dimensional contractible simplicial complex. Hence, it is of type* F. *The group* PSOut_n *has the same properties, with* $n-1$ *replaced by* $n-2$.

Proof Collins [Col89] proved that the groups SAut_n and SOut_n admit proper simplicial cocompact actions on contractible simplicial complexes of dimensions $n-1$ and $n-2$, respectively. On the other hand, PSAut_n and PSOut_n are torsion-free by a result of Baumslag and Taylor [BT68]. Since a proper action of a torsion-free group must be free, restricting the above actions of SAut_n and SOut_n to their finite index subgroups PSAut_n and PSOut_n, we obtain free actions which remain cocompact. ∎

2.3 Setting the stage for BNSR-invariants of McCool groups

We can already begin to see some facts relevant to the BNSR-invariants by looking at the McCool presentation. First, we can compute abelianizations. Since each defining relation is a commutator, the abelianization of PSAut_n is freely generated by the images of the α_{ij} and thus is isomorphic to $\mathbb{Z}^{n(n-1)}$. Hence, $\mathrm{Hom}(\mathrm{PSAut}_n, \mathbb{R}) \cong \mathbb{R}^{n(n-1)}$, and so $\mathbb{S}(\mathrm{PSAut}_n) \cong S^{n(n-1)-1}$. As for PSOut_n, in the abelianization, we take $\mathbb{Z}^{n(n-1)}$

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

generated by the images $\overline{\alpha}_{ij}$ of the α_{ij}, and we mod out n independent relations, of the form $\overline{\alpha}_{[n]\setminus\{j\},j} = 0$ for each j. Thus, the abelianization of PSOut_n is isomorphic to $\mathbb{Z}^{n(n-1)-n} = \mathbb{Z}^{n(n-2)}$, so $\mathbb{S}(\mathrm{PSOut}_n) \cong S^{n(n-2)-1}$.

Lemma 2.7 *The BNSR-invariants of* PSAut_n *and* PSOut_n *are symmetric.*

Proof For PSAut_n this is [Zar18, Observation 2.3], and the argument for PSOut_n is analogous. To spell it out, from the McCool presentation, we see that there is an automorphism $\omega \in \mathrm{Aut}(\mathrm{PSAut}_n)$ sending each α_{ij} to α_{ij}^{-1}, which satisfies $\chi \circ \omega = -\chi$ for every character χ of PSAut_n, and the induced automorphism $\overline{\omega}$ of PSOut_n has the same property. ■

As we already mentioned, in [OK00], Orlandi-Korner fully computed the first BNSR-invariant $\Sigma^1(\mathrm{PSAut}_n)$ for all n. In [Zar18], the second author gave a partial computation of $\Sigma^m(\mathrm{PSAut}_n)$ for $m \geq 2$, and in particular, proved that all "positive" (and all "negative") characters of PSAut_n lie in $\Sigma^{n-2}(\mathrm{PSAut}_n)\setminus\Sigma^{n-1}(\mathrm{PSAut}_n)$.

Remark 2.8 One can also ask about the BNSR-invariants of the full (as opposed to pure) symmetric automorphism groups SAut_n and SOut_n, but this turns out to be a much easier question since these groups have very low-dimensional character spheres. We have $\mathrm{Hom}(\mathrm{SAut}_n, \mathbb{R}) \cong \mathbb{R}$, so $\mathbb{S}(\mathrm{SAut}_n) = S^0$, and $\Sigma^k(\mathrm{SAut}_n)$ is all of $\mathbb{S}(\mathrm{SAut}_n)$ for $k \leq n-2$ and empty for $k \geq n-1$ [Zar18, Theorem B]. Even more extreme, $\mathrm{Hom}(\mathrm{SOut}_n, \mathbb{R})$ is trivial, since any nontrivial map $\mathrm{SAut}_n \to \mathbb{R}$ must restrict nontrivially to $\mathrm{Inn}(F_n)$.

3 Empty BNSR-invariants

In this section, we prove parts (b) and (d) of Theorem 1.1, which assert that certain invariants are empty. This is the easier part; indeed, it was arguably already known to experts, although some pieces of proof need to be assembled.

Let G be a group such that all the rational homology groups $H_i(G, \mathbb{Q})$ of G are finite-dimensional and vanish for large enough i. For example, this happens if G is of type FP, meaning the trivial G-module $\mathbb{Z}$ admits a finite-length resolution by finitely generated projective G-modules. In particular, this happens if G is of type F. The *Euler characteristic* $\chi(G)$ of G is

$$\chi(G) = \sum_{i=0}^{\infty} (-1)^i \dim H_i(G, \mathbb{Q}).$$

The following result must be well known, but we could not find it stated in this concise form in the literature; the proof is reasonably elementary.

Proposition 3.1 *Let G be a group of type* F *such that $\chi(G) \neq 0$. Then $\Sigma^\infty(G) \cap -\Sigma^\infty(G) = \varnothing$, so in particular, if the BNSR-invariants of G are symmetric, then $\Sigma^\infty(G) = \varnothing$.*

Proof Suppose $\Sigma^\infty(G) \cap -\Sigma^\infty(G) \neq \varnothing$. Since G is of type F, $\Sigma^\infty(G) = \Sigma^m(G)$, where $m < \infty$ is the dimension of some finite classifying space for G, so $\Sigma^m(G) \cap -\Sigma^m(G) \neq \varnothing$. Since $\Sigma^m(G)$ is open in $\mathbb{S}(G)$, we can choose a discrete character $0 \neq \chi\colon G \to \mathbb{Z}$ such that $[\pm\chi] \in \Sigma^m(G) = \Sigma^\infty(G)$. Let $K = \ker(\chi)$, so K is of type F_∞.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

Since G has a finite classifying space of dimension m, its cohomological dimension $cd(G)$ is at most m and hence $cd(K) \leq cd(G) < \infty$ (see, e.g., [Bro82, Propositions VIII.2.2 and VIII.2.4]). Combined with K being of type F_∞, hence FP_∞, this implies that K is of type FP by [Bro82, Proposition VIII.6.1]. Now, by [Bro82, Proposition IX.7.3(d)], since K, G, and $\mathbb{Z}$ are all of type FP, the Euler characteristic $\mathcal{X}$ is multiplicative for the short exact sequence $1 \to K \to G \to \mathbb{Z} \to 1$, and we get $\mathcal{X}(G) = \mathcal{X}(K) \cdot \mathcal{X}(\mathbb{Z}) = \mathcal{X}(K) \cdot 0 = 0$, a contradiction. ∎

Remark 3.2 The above proof applies equally well to the homological BNSR-invariant $\Sigma^\infty(G;\mathbb{Z})$, since we only need to know that K was of type FP_∞ (not necessarily F_∞). We should also mention that, as explained to us by Dawid Kielak, using some deep results from ℓ^2-homology the assumption about the BNSR-invariants being symmetric can be relaxed, and the conclusion can be strengthened to saying that $\Sigma^\infty(G;\mathbb{Q}) = \varnothing$, where $\Sigma^\infty(G;\mathbb{Q})$ is the homological BNSR-invariants with $\mathbb{Q}$ coefficients. This follows from Theorem A of [HK24], together with the fact that for groups of type F the Euler characteristic and ℓ^2-Euler characteristic coincide (see [Lüc02, Theorem 6.80(1)]). For our purposes, we only need the result as stated in Proposition 3.1.

We can now prove Theorem 1.1(b)(d), asserting that $\Sigma^{n-1}(\mathrm{PSAut}_n) = \varnothing$ and $\Sigma^{n-2}(\mathrm{PSOut}_n) = \varnothing$.

Proof of Theorem 1.1(b)(d) By Proposition 2.6, PSAut_n and PSOut_n are of type F. By Lemma 2.7, the BNSR-invariants of these groups are symmetric. Finally, it is known that $\mathcal{X}(\mathrm{PSAut}_n) \neq 0$ and $\mathcal{X}(\mathrm{PSOut}_n) \neq 0$; in fact, $\mathcal{X}(\mathrm{PSAut}_n) = (1-n)^{n-1} \neq 0$ and $\mathcal{X}(\mathrm{PSOut}_n) = (1-n)^{n-2} \neq 0$ (see Theorems 1.2 and 8.2 of [MM04]). Thus, $\Sigma^{n-1}(\mathrm{PSAut}_n) = \varnothing$ and $\Sigma^{n-2}(\mathrm{PSOut}_n) = \varnothing$ by Proposition 3.1. ∎

4 Dense BNSR-invariants

In this section, we prove the much more difficult, "positive" statements of Theorem 1.1, parts (a) and (c), asserting that $\Sigma^{n-2}(\mathrm{PSAut}_n)$ and $\Sigma^{n-3}(\mathrm{PSOut}_n)$ are dense in their respective character spheres. The general plan is to apply a criterion of Meier, Meinert, and Van Wyk, with the bulk of the work relating to the Whitehead poset WO_n and McCullough–Miller space MMO_n of PSOut_n, along with natural "auter" variants. In the coming sections, we formulate the Meier–Meinert–Van Wyk criterion, discuss the Whitehead poset and McCullough–Miller space, and then prove the density results.

4.1 Higher generation and the Meier–Meinert–Van Wyk criterion

The key to proving that certain characters are in certain BNSR-invariants will be a general criterion due to Meier, Meinert, and Van Wyk [MMV98] (see Theorem 4.3 below), involving the concept of higher generation due to Abels and Holz. We start with several definitions.

Let G be a group, let $\mathcal{H}$ be a family of subgroups of G, and let $\mathrm{Cos}(G,\mathcal{H})$ be the collection of all cosets gH with $g \in G$ and $H \in \mathcal{H}$. We will view both $\mathcal{H}$ and $\mathrm{Cos}(G,\mathcal{H})$ as posets with respect to inclusion. Recall that given a poset P, the associated flag complex $\mathcal{F}(P)$ is the simplicial complex whose simplices are finite chains of elements

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

of P (this is also known as the order complex). If P has a global minimum element $\widehat{0}$, let us also write

$$\mathcal{F}_0(P) = \mathcal{F}(P\backslash\{\widehat{0}\}).$$

Below, we will be concerned with two flag complexes associated with G and $\mathcal{H}$. The first is $\mathcal{F}(\mathrm{Cos}(G,\mathcal{H}))$, the flag complex of the poset of all cosets of subgroups in $\mathcal{H}$. The second is $\mathcal{F}_0(\mathcal{H})$, the flag complex of the poset of all nontrivial subgroups in $\mathcal{H}$ (only the nontrivial subgroups, not all cosets thereof).

Let $\mathcal{U} = \{U_i \mid i \in I\}$ be a covering of a set X. Recall that the *nerve* of $\mathcal{U}$, denoted by $N(\mathcal{U})$, is the simplicial complex with vertex set I such that $i_0, \dots, i_p$ span a p-simplex whenever $U_{i_0} \cap \cdots \cap U_{i_p}$ is non-empty.

The following definition was introduced by Abels and Holz in [AH93].

Definition 4.1 (n-generating) Let $\mathcal{H}$ be a family of subgroups of a group G, so $\mathrm{Cos}(G,\mathcal{H})$ is a covering of G. For $n \in \mathbb{N}$, the family $\mathcal{H}$ is said to be *n-generating* for G if the nerve $N(\mathrm{Cos}(G,\mathcal{H}))$ is $(n-1)$-connected, and *∞-generating* if the nerve is contractible.

Remark 4.2 If $\mathcal{H}$ is closed under finite intersections, then by [AH93, Theorem 1.4(b)], $N(\mathrm{Cos}(G,\mathcal{H}))$ is homotopy equivalent to the flag complex $\mathcal{F}(\mathrm{Cos}(G,\mathcal{H}))$. Thus, in this case, we can replace $N(\mathrm{Cos}(G,\mathcal{H}))$ in Definition 4.1 with the complex $\mathcal{F}(\mathrm{Cos}(G,\mathcal{H}))$, which is somewhat easier to understand, being the flag complex of a poset.

We are now ready to state the general result of Meier–Meinert–Van Wyk from [MMV98] relating BNSR-invariants to highly generating families of subgroups.

Theorem 4.3 [MMV98, Theorem 3.1] *Let G be a group of type F_n. Let $\mathcal{H}$ be a non-empty finite intersection-closed n-generating family of subgroups of G such that each $H \in \mathcal{H}$ is of type F_n. Let $\chi\colon G \to \mathbb{R}$ be a character such that $\chi|_H \neq 0$ for each nontrivial subgroup $H \in \mathcal{H}$. Suppose that the following hold:*

(i) *Either the trivial subgroup does not belong to $\mathcal{H}$ or the simplicial complex $\mathcal{F}_0(\mathcal{H})$ (as defined above) is $(n-1)$-connected, and*
(ii) *$[\chi|_H] \in \Sigma^n(H)$ for all nontrivial $H \in \mathcal{H}$.*

Then $[\chi] \in \Sigma^n(G)$.

Remark 4.4 The condition that the trivial subgroup does not belong to $\mathcal{H}$ is a strong constraint since $\mathcal{H}$ needs to be closed under intersections. On the other hand, if the trivial subgroup does belong to $\mathcal{H}$, the complex $\mathcal{F}(\mathcal{H})$ is trivially contractible, which explains why we are looking at $\mathcal{F}_0(\mathcal{H})$ rather than $\mathcal{F}(\mathcal{H})$ in this case.

4.2 The Whitehead poset

In this section, we recall the definition and prove some properties of the Whitehead poset WO_n, generally following [BMMM01]. The elements of WO_n can be viewed combinatorially as bipartite labeled trees on $[n] = \{1, 2, \dots, n\}$, which we now define.

Definition 4.5 A *bipartite labeled tree on [n]* is a tree that has n vertices labeled by the integers from 1 to n and some finite number of unlabeled vertices, satisfying the following conditions:

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press



Figure 1: Bipartite labeled trees on $[4]$, with the second tree obtained from the first tree by a folding at the vertex labeled 3.

(i) each edge connects a labeled vertex to an unlabeled vertex;
(ii) each unlabeled vertex has degree at least 2.

The set of all bipartite labeled trees on $[n]$ will be denoted by WO_n. We will now define a partial order on WO_n. If $T \in \mathrm{WO}_n$ has a labeled vertex v connected by edges to two unlabeled vertices w_1 and w_2, we can form a new tree T' by identifying w_1 and w_2 and identifying the edges joining w_1 and w_2 to v. In this case, we say that T' is obtained from T by a *folding* at v. See Figure 1 for an example.

Define a partial order $\preceq$ on WO_n as follows: given $T_1, T_2 \in \mathrm{WO}_n$, we set $T_1 \preceq T_2$ if T_1 can be obtained from T_2 by a (possibly empty) sequence of foldings.

Definition 4.6 The set WO_n with the partial order $\preceq$ is called the *Whitehead poset*.

We will also be working with WA_n, the "auter" version of the Whitehead poset. We define WA_n to be the subposet of WO_{n+1} consisting of all trees T for which the vertex labeled $n+1$ has degree 1.

If $T \in \mathrm{WO}_n$ has k unlabeled vertices, define the *degree* of T to be

$$\deg(T) = k - 1.$$

Let us collect a few basic properties of foldings and the degree function on WO_n.

Observation 4.7 *The following hold:*

(a) *If T' is obtained from T by a folding, then* $\deg(T') = \deg(T) - 1$.
(b) $0 \le \deg(T) \le n-2$ *for all* $T \in \mathrm{WO}_n$.
(c) *Every tree of positive degree admits at least one folding.*
(d) *Every tree of degree* $< n-2$ *admits at least one unfolding (the reverse of a folding).*
(e) WO_n *contains a unique tree of degree* 0 *(in which the unique unlabeled vertex is connected to all n labeled vertices), and this tree is the (unique) global minimum of* WO_n.
(f) *The maximal elements of* WO_n *are precisely the trees of degree* $n-2$. *Moreover, for every maximal tree T, there exists a chain* $T_0 \prec T_1 \prec \cdots \prec T_{n-2} = T$ *in* WO_n.
(g) *The maximal elements of* WA_{n-1} *are precisely the trees of degree* $n-2$ *which lie in* WA_{n-1} *(recall that* WA_{n-1} *is defined as a subset of* WO_n*). Moreover, for every maximal tree* $T \in WA_{n-1}$, *every chain as in (f) is contained in* WA_{n-1}.

Proof (a) This is clear from the definition of folding.

(b) Suppose that T has k unlabeled vertices (i.e., $k = \deg(T) + 1$). Since every edge connects an unlabeled vertex to a labeled vertex, we must have $k \ge 1$, so $\deg(T) \ge 0$. The total number of vertices in T is $n + k$ and hence (since T is a tree), it has exactly

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

$n+k-1$ edges. Since in addition, every unlabeled vertex has degree ≥ 2, we have $n+k-1 \geq 2k$, that is, $k \leq n-1$, or equivalently $\deg(T) \leq n-2$, which proves (b).

(c) A tree admits no foldings if and only if all labeled vertices have degree 1. This happens exactly when $n+k-1$ (the total number of edges) equals n (the total number of labeled vertices), that is, when $k=1$ or, equivalently, $\deg(T)=0$.

(d) A tree T admits an unfolding if and only if it has an unlabeled vertex of degree >2, and by the computation in (b), the latter happens exactly when $\deg(T) < n-2$.

(e) The first assertion is clear and the second one follows from (c).

(f) If $\deg(T) = n-2$, then T is maximal by (a) and (b), and if $\deg(T) < n-2$, then T is not maximal by (d). The last assertion of (f) follows from (c) and (a).

(g) As in (f), if $\deg(T) = n-2$, then T is maximal. If $\deg(T) < n-2$, then by the proof of (d) T has an unlabeled vertex of degree >2, call it v. At least one of the neighbors of v is labeled by some $i \neq n+1$, and an unfolding at that neighbor produces a tree $T' \in \mathrm{WA}_{n-1}$ with $T \prec T'$, so T is not maximal. The last assertion of (g) holds simply because foldings do not increase the degrees of labeled vertices, so if $T \in \mathrm{WA}_{n-1}$, then $S \in \mathrm{WA}_{n-1}$ for any $S \prec T$. ∎

The following result is an immediate consequence of parts (f) and (g) of Observation 4.7.

Corollary 4.8 *The flag complex $\mathcal{F}(\mathrm{WO}_n)$ is $(n-2)$-dimensional, and the flag complex $\mathcal{F}(\mathrm{WA}_n)$ is $(n-1)$-dimensional.*

4.3 A family of abelian subgroups

The main goal of this section is to show that the Whitehead poset is isomorphic to a certain poset of abelian subgroups of PSOut_n ordered by inclusion. To build up to this, let us construct a subgroup $H(T) \leq \mathrm{PSOut}_n$ associated with each $T \in \mathrm{WO}_n$. Given $T \in \mathrm{WO}_n$ and $j \in [n]$, let T_j be the graph obtained from T by removing the vertex labeled j (and all edges incident to it), and consider the partition of $[n]\backslash\{j\}$ where two elements lie in the same block if and only if the vertices of T_j with the corresponding labels lie in the same connected component. Denote this partition by $\mathcal{P}(T,j)$. Now let

$$H(T) = \left\langle \alpha_{I,j} \;\middle|\; j \in [n],\, I \in \mathcal{P}(T,j) \right\rangle / \mathrm{Inn}(F_n),$$

and write

$$\mathcal{HO}_n = \{H(T) \mid T \in \mathrm{WO}_n\}.$$

Proposition 4.9 *The group $H(T)$ is free abelian of rank* $\deg(T)$.

Proposition 4.9 is easy to deduce from [MM96, Proposition 5.1] (see also [JMM06, Theorem 2.3(ii)]); however, since the groups $H(T)$ are introduced in [MM96] in a different way and moreover [MM96] considers symmetric automorphisms of arbitrary free products (as opposed to free groups in our article) making the statements more technical, we will provide a direct proof for the convenience of the reader.

Below by $\overline{\alpha}_{i,j}$, we denote the image of $\alpha_{i,j}$ in PSOut_n. We first establish an auxiliary result.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

Lemma 4.10 *For each $1 \le j \le n$, let $G_j = \langle \overline{\alpha}_{i,j} \mid i \ne j \rangle$. The following hold:*

(a) *The group G_j is abelian.*

(b) *Fix $1 \le j \le n$. For $c_i \in \mathbb{Z}$, the product $\prod_{j \ne i} \overline{\alpha}_{i,j}^{c_i}$ equals 1 if and only if all c_i are the same. Hence, $G_j \cong \mathbb{Z}^{n-2}$.*

(c) *Let $g = \prod_{j=1}^{n} g_j$ with $g_j \in G_j$ for each j. Then $g = 1$ if and only if each $g_j = 1$.*

Proof (a) This is immediate from Lemma 2.5(i).

(b) The "if" part is clear. The "only if" part holds since an automorphism of the form $\prod_{j \ne i} \alpha_{i,j}^{c_i}$ always fixes x_j, so it can only be inner if it is the conjugation by a power of x_j.

(c) Let $V = \mathrm{PSOut}_n / [\mathrm{PSOut}_n, \mathrm{PSOut}_n]$, and let $\pi : \mathrm{PSOut}_n \to V$ be the natural projection. As we already explained in Section 2.3, $V \cong \mathbb{Z}^{n(n-2)}$. Since $V = \sum_{j=1}^{n} \pi(G_j)$ and each $G_j \cong \mathbb{Z}^{n-2}$ by (b), it follows that π is injective on each G_j and $V = \oplus_{j=1}^{n} \pi(G_j)$. This implies (c). ∎

Proof of Proposition 4.9 We first prove that $H(T)$ is abelian. Take any two generators $\alpha_{I,j}$ and $\alpha_{K,m}$ of $H(T)$, so $I \in \mathcal{P}(T, j)$ and $K \in \mathcal{P}(T, m)$. If $j = m$, these generators commute by Lemma 2.5(ii).

Suppose now that $j \ne m$. Since $\overline{\alpha}_{I,j}^{-1} = \overline{\alpha}_{[n]\setminus(I\cup\{j\}),j}$, replacing $\alpha_{I,j}$ by $\alpha_{[n]\setminus(I\cup\{j\}),j}$ if needed, we can assume that $m \notin I$ (and hence $m \notin I \cup \{j\}$). Since $I \in \mathcal{P}(T, j)$, by definition of the partition $\mathcal{P}(T, j)$, there exists a subtree of T whose set of vertex labels is precisely $I \cup \{j\}$. Since $m \notin I \cup \{j\}$, the set $I \cup \{j\}$ must lie in a single block of the partition $\mathcal{P}(T, m)$. Hence, either $I \cup \{j\} \subseteq K$ or $I \cup \{j\}$ is disjoint from K (and hence from $K \cup \{m\}$). In each case, $\alpha_{I,j}$ and $\alpha_{K,m}$ commute by Lemma 2.5(iii).

Thus, we have proved that $H(T)$ is abelian, so any element of $H(T)$ can be written as a product $\prod_{j=1}^{n} p_j$, where each $p_j = \prod_{I \in \mathcal{P}(T,j)} \overline{\alpha}_{I,j}^{c_{I,j}}$ for some $c_{I,j} \in \mathbb{Z}$. By Lemma 4.10(c), $\prod_{j=1}^{n} p_j = 1$ if and only if $p_j = 1$ for each $1 \le j \le n$. For a fixed j, the elements of $\mathcal{P}(T, j)$ are pairwise disjoint, so by Lemmas 4.10(b) and 2.5(i), we have $p_j = 1$ if and only if the integers $c_{I,j}$ are the same for all I. Hence, $H(T)$ is free abelian of rank $\sum_{j=1}^{n} (d_j - 1) = \sum_{j=1}^{n} d_j - n$, where d_j is the degree of the vertex of T labeled j.

Since every edge connects a labeled vertex to an unlabeled vertex, $\sum_{j=1}^{n} d_j$ is the total number of edges of T, which is also equal to $n + \deg(T)$ since T has n labeled vertices and $\deg(T) + 1$ unlabeled vertices. Thus, the rank of $H(T)$ is equal to $(n + \deg(T)) - n = \deg(T)$, as desired. ∎

In the "auter" case, for each $T \in \mathrm{WA}_n$, let

$$H_A(T) = H(T) \cap \mathrm{PSAut}_n,$$

where we view PSAut_n as a subgroup of PSOut_{n+1} by having each automorphism act trivially on x_{n+1} (and noting that nontrivial automorphisms of this form are never

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

inner). Let

$$\mathcal{HA}_n = \{H_A(T) \mid T \in \mathrm{WA}_n\}.$$

Note that each $H_A(T)$ is abelian, being a subgroup of $H(T)$.

We will show that the map $T \mapsto H(T)$ is an isomorphism of posets between WO_n and $\mathcal{HO}_n$ (ordered by inclusion) and likewise the map $T \mapsto H_A(T)$ is an isomorphism of posets between WA_n and $\mathcal{HA}_n$. First, we need some preparations.

Definition 4.11 Let $T \in \mathrm{WO}_n$ and $j \in [n]$. A subset I of $[n]\backslash\{j\}$ will be called *(T,j)-complete* if I is a (possibly empty) union of some blocks of the partition $\mathcal{P}(T, j)$.

Lemma 4.12 *Let $T \in WO_n$ and $I \subseteq [n]\backslash\{j\}$. Then $\overline{\alpha}_{I,j} \in H(T)$ if and only if I is (T, j)-complete.*

Proof If I is (T, j)-complete, then $I = \sqcup_k I_k$, where each $I_k \in \mathcal{P}(T, j)$ and hence $\overline{\alpha}_{I,j} = \prod_k \overline{\alpha}_{I_k,j} \in H(T)$.

Let us now prove the converse. Suppose that $\overline{\alpha}_{I,j} \in H(T)$ for some $j \in [n]$ and $I \subseteq [n]\backslash\{j\}$, so that

$$\overline{\alpha}_{I,j} = \prod_{t=1}^{k} (\overline{\alpha}_{B_t,j_t})^{c_t},$$

where $B_t \in \mathcal{P}(T, j_t)$ for each t, all pairs (B_t, j_t) are distinct, and all $c_t \neq 0$. By Lemma 4.10(c), we can assume that $j_t = j$ for each t in the above equality, and we can rewrite the latter as

$$(***) \qquad \prod_{i \in I} \overline{\alpha}_{i,j} = \prod_{t=1}^{k} \left(\prod_{i \in B_t} \overline{\alpha}_{i,j} \right)^{c_t}.$$

Suppose now that some $i, i' \in [n]\backslash\{j\}$ lie in the same block of the partition $\mathcal{P}(T, j)$. Then $\overline{\alpha}_{i,j}$ and $\overline{\alpha}_{i',j}$ appear on the right-hand side of (***) with the same exponent. Since the only vanishing products of the elements $\overline{\alpha}_{i,j}$ for a fixed j are powers of $\prod_{i \in [n]\backslash\{j\}} \overline{\alpha}_{i,j}$, it follows that $\overline{\alpha}_{i,j}$ and $\overline{\alpha}_{i',j}$ also have the same exponent on the left-hand side of (***), so i and i' either both lie in I or both lie outside of I. It follows that I is a union of blocks of $\mathcal{P}(T, j)$, so it is (T, j)-complete. ∎

Lemma 4.13 *The function $\eta : T \mapsto H(T)$ is a bijection from WO_n to $\mathcal{HO}_n$, and the function $\eta_A : T \mapsto H_A(T)$ is a bijection from WA_n to $\mathcal{HA}_n$.*

Proof The functions η and η_A are surjective by construction, so we just need to prove injectivity.

We first prove that η is injective by describing an algorithm for extracting T from $H(T)$. Fix $T \in \mathrm{WO}_n$. We claim that for any distinct $i, j \in [n]$, the following conditions are equivalent:

(i) The vertices labeled i and j are connected to the same unlabeled vertex in T.
(ii) The unique reduced path in T connecting the vertices labeled i and j contains no other labeled vertices.
(iii) The labels i and j lie in the same block of the partition $\mathcal{P}(T, k)$ for any $k \in [n]\backslash\{i, j\}$.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

(iv) For any $\overline{\alpha}_{I,k} \in H(T)$ with $k \neq i, j$, the set I contains either both i and j or neither of them.

The equivalences (i) $\Longleftrightarrow$ (ii) and (ii) $\Longleftrightarrow$ (iii) are clear and (iii) $\Longleftrightarrow$ (iv) holds by Lemma 4.12. Since condition (iv) can be read off from $H(T)$, the same is true for (i).

Let us now form an auxiliary graph $\mathrm{Aux}(T)$. It has n vertices, labeled by the elements of $[n]$, and no unlabeled vertices; the vertices labeled by i and j are connected by an edge if and only if the vertices of T labeled i and j are connected by edges to the same unlabeled vertex in T. By the previous paragraph, $\mathrm{Aux}(T)$ is determined by $H(T)$. On the other hand, it is easy to recover T from $\mathrm{Aux}(T)$: create an unlabeled vertex for each maximal clique C in $\mathrm{Aux}(T)$ and connect that vertex to a vertex labeled by i for each $i \in C$. Thus, we have proved that η is injective.

We now deduce injectivity of η_A from that of η. For this, we need to prove that if $T, U \in \mathrm{WA}_n$ are such that $H(T) \cap \mathrm{PSAut}_n = H(U) \cap \mathrm{PSAut}_n$ then $H(T) = H(U)$. By symmetry, it suffices to check that $H(T) \subseteq H(U)$. Consider any $\overline{\alpha}_{I,j} \in H(T)$, so $j \in [n+1]$ and I is (T, j)-complete. We will prove that $\overline{\alpha}_{I,j} \in H(U)$ by splitting into three cases.

Case 1: Suppose $j \in [n]$ and $n+1 \notin I$. Then $\overline{\alpha}_{I,j} \in \mathrm{PSAut}_n$, hence the hypothesis $H(T) \cap \mathrm{PSAut}_n = H(U) \cap \mathrm{PSAut}_n$ ensures that $\overline{\alpha}_{I,j} \in H(U)$.

Case 2: Suppose $j = n+1$. Since $n+1$ is a leaf in T, we have $I = [n]$ or $\varnothing$; in either case, $\alpha_{I,j}$ is inner, so $\overline{\alpha}_{I,j} = 1$ is automatically in $H(U)$.

Case 3: Suppose $j \in [n]$ and $n+1 \in I$. Let I' be the complement of I in $[n+1]\backslash\{j\}$, so $\alpha_{I,j}\alpha_{I',j}$ is inner. Hence, $\overline{\alpha}_{I',j} = \overline{\alpha}_{I,j}^{-1} \in H(T)$, but now $n+1 \notin I'$, so by case 1, we have $\overline{\alpha}_{I',j} \in H(U)$, and thus $\overline{\alpha}_{I,j} \in H(U)$ as desired. ∎

Recall that a poset is a *meet semilattice* if any two elements x and y have a unique greatest lower bound, called their *meet*, denoted $x \wedge y$.

Proposition 4.14 *[MM96, Lemmas 5.4 and 5.5] The posets* WO_n *and* WA_n *are meet semilattices. Moreover,* $H(T \wedge U) = H(T) \cap H(U)$ *for any* $T, U \in \mathrm{WO}_n$ *and* $H_A(T \wedge U) = H_A(T) \cap H_A(U)$ *for any* $T, U \in \mathrm{WA}_n$*. In particular, the families* $\mathcal{HA}_n$ *and* $\mathcal{HO}_n$ *are closed under taking intersections and so are also meet semilattices.*

Now let us prove that the bijection $\eta : T \mapsto H(T)$ is a poset isomorphism from WO_n to $\mathcal{HO}_n$, along with the "auter" version.

Proposition 4.15 *For* $T, U \in \mathrm{WO}_n$*, we have* $T \preceq U$ *if and only if* $H(T) \leq H(U)$*. For* $T, U \in \mathrm{WA}_n$*, we have* $T \preceq U$ *if and only if* $H_A(T) \leq H_A(U)$*. Hence,* η *and* η_A *are isomorphisms of posets, and we get* $\mathcal{F}(\mathrm{WO}_n) \cong \mathcal{F}(\mathcal{HO}_n)$ *and* $\mathcal{F}(\mathrm{WA}_n) \cong \mathcal{F}(\mathcal{HA}_n)$*. Analogous results also hold with* $\mathcal{F}$ *replaced by* $\mathcal{F}_0$ *everywhere.*

Proof Suppose that $T \preceq U$, so T is obtained from U by a sequence of foldings. Then for any $j \in [n]$ and any $I \subseteq [n]\backslash\{j\}$ that is (T, j)-complete, I is also (U, j)-complete. We conclude that $H(T) \leq H(U)$.

Now suppose that $H(T) \leq H(U)$. This implies that $H(T) = H(T) \cap H(U)$, and so by Proposition 4.14, we have $H(T) = H(T \wedge U)$. By Lemma 4.13, this tells us that $T = T \wedge U$, and so $T \preceq U$. Since we already know that $\eta : \mathrm{WO}_n \to \mathcal{HO}_n$ is a bijection, this proves that it is an isomorphism of posets.

The proof for η_A is analogous, and the remaining statements in the proposition are all immediate consequences of the facts that η and η_A are poset isomorphisms. ∎

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

4.4 McCullough–Miller space

We are now ready to define the McCullough–Miller spaces MMO_n and MMA_n, slightly rephrasing the original definitions from [CGJ05, MM96]. The idea for MMO_n is to glue together copies of $\mathcal{F}(\mathrm{WO}_n)$, one for each element of PSOut_n, along appropriate intersections. Since we have already established that WO_n is isomorphic to $\mathcal{HO}_n$, it is not surprising that the end result of this is a complex isomorphic to $\mathcal{F}(\mathrm{Cos}(\mathrm{PSOut}_n, \mathcal{HO}_n))$, and indeed we will simply define McCullough–Miller space this way. The equivalence between this viewpoint and the original viewpoint of McCullough–Miller in [MM96] is not difficult, and is discussed, for example, in [Gri13] (along with the MMA_n case).

Definition 4.16 (McCullough–Miller space) Define the *McCullough–Miller space* MMO_n to be

$$\mathrm{MMO}_n = \mathcal{F}(\mathrm{Cos}(\mathrm{PSOut}_n, \mathcal{HO}_n)),$$

and define the "auter" version MMA_n to be

$$\mathrm{MMA}_n = \mathcal{F}(\mathrm{Cos}(\mathrm{PSAut}_n, \mathcal{HA}_n)).$$

The only property of the McCullough–Miller spaces that we will need is also the most important one, namely, that they are contractible.

Proposition 4.17 *The complexes* MMO_n *and* MMA_n *are contractible. Thus,* $\mathcal{HO}_n$ *and* $\mathcal{HA}_n$ *are* ∞*-generating, for* PSOut_n *and* PSAut_n *respectively.*

Proof Contractibility of MMO_n was proved by McCullough and Miller in [MM96, Section 4]. Contractibility of MMA_n is essentially proved by Chen, Glover, and Jensen in [CGJ05, Theorem 3.1]. Technically, in [CGJ05], Chen, Glover, and Jensen consider pure symmetric automorphisms not of free groups, but of free products of finite groups. However, the proof of contractibility of the associated complex works for free products of any groups, in particular free groups, as explained by Griffin in [Gri13, Section 5] (see [Gri13, Theorem 5.5]). The statement about ∞-generation is now immediate from Remark 4.2. ∎

4.5 Proof of density

In this section, we prove parts (a) and (c) of Theorem 1.1, asserting that $\Sigma^{n-2}(\mathrm{PSAut}_n)$ and $\Sigma^{n-3}(\mathrm{PSOut}_n)$ are dense in their respective character spheres. The idea is to prove that a suitably generic character satisfies the hypotheses of Theorem 4.3, using the family of subgroups $\mathcal{HA}_n$ for PSAut_n and $\mathcal{HO}_n$ for PSOut_n.

Most of the hypotheses involved in Theorem 4.3 have been dealt with by now. The one we still need to prove is that the flag complexes of nontrivial subgroups from the above families are highly connected. We will actually prove a stronger statement, involving the property of being homotopy Cohen–Macaulay.

Definition 4.18 (Homotopy Cohen–Macaulay) A simplicial complex X of finite dimension n is called *homotopy Cohen–Macaulay* if for every p-simplex σ, the link $\mathrm{lk}(\sigma)$ is $(n-p-2)$-connected. (This includes the empty simplex, so X itself

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

should be $(n-1)$-connected.) A poset is called *homotopy Cohen–Macaulay* if its flag complex is.

Proposition 4.19 *The posets* WO_n *and* WA_n *are homotopy Cohen–Macaulay. Hence, the flag complex* $\mathcal{F}_0(\mathcal{HO}_n)$ *is* $(n-4)$*-connected and* $\mathcal{F}_0(\mathcal{HA}_n)$ *is* $(n-3)$*-connected.*

Remark 4.20 Since Proposition 4.8 says that the flag complexes of WO_n and WA_n are $(n-2)$-dimensional and $(n-1)$-dimensional, respectively, and since $\mathcal{F}_0(\mathcal{HO}_n)$ and $\mathcal{F}_0(\mathcal{HA}_n)$ are isomorphic to links of 0-simplices in the flag complexes of WO_n and WA_n, respectively, namely, the trivial vertices, the second statement follows from the first.

The fact that WO_n is homotopy Cohen–Macaulay has already been proved in [BMMM01, Theorem 5.13]. To show that WA_n is also homotopy Cohen–Macaulay, we need to recreate and adjust the proof for WO_n. First, we will introduce some terminology dealing with posets.

Definition 4.21 A poset P is called

- *bounded* if it has a global minimum and a global maximum;
- *graded* if it is bounded and all maximal chains have the same length, which is then called the *length* of P.

Given $p \leq q$ in a poset P, write $[p,q] = \{r \mid p \leq r \leq q\}$, and call this a (closed) *interval*. Given $p, q \in P$ with $p \neq q$, we say that q *covers* p if $[p,q] = \{p,q\}$. An *atom* in a poset with a global minimum is any element that covers the global minimum.

There is a nice sufficient condition (see [BMMM01, Lemma 5.3]), originally due to Björner and Wachs [BW83], for a bounded graded poset P to be homotopy Cohen–Macaulay, namely, that it admits a recursive atom ordering.

Definition 4.22 Let P be a bounded graded poset with global minimum $\widehat{0}$ and global maximum $\widehat{1}$. We say that P *admits a recursive atom ordering* if either P has length 1 or P has length greater than 1 and there is an ordering $a_1, \ldots, a_\ell$ of the atoms of P such that the following hold:

(i) For each $1 \leq j \leq \ell$, the interval $[a_j, \widehat{1}]$ admits a recursive atom ordering such that any atoms of $[a_j, \widehat{1}]$ that cover some a_i with $i < j$ come first in the ordering.
(ii) For all $i < j$, if $a_i, a_j < y$, then there exist $i' < j$ and $z \leq y$ such that z covers $a_{i'}$ and a_j.

Let $\widehat{\mathrm{WO}}_n$ be the result of adding a new element to WO_n, as a global maximum. Construct $\widehat{\mathrm{WA}}_n$ from WA_n in the same way. Let ZO_n and ZA_n be the dual posets of $\widehat{\mathrm{WO}}_n$ and $\widehat{\mathrm{WA}}_n$, respectively, that is, replace every $\leq$ with $\geq$. By [BMMM01, Theorem 5.12], ZO_n admits a recursive atom ordering. Our goal is to adapt that proof to the case of ZA_n.

Proposition 4.23 ZA_n *admits a recursive atom ordering.*

First, let us explain why Proposition 4.19 for WA_n follows from Proposition 4.23.

Proof of Proposition 4.19 for WA_n Proposition 4.23 and [BMMM01, Lemma 5.3] imply that ZA_n is homotopy Cohen–Macaulay. It is immediate from the definitions

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

that the flag complex of WA_n is isomorphic to the link of the global minimum in ZA_n, and that the property of being homotopy Cohen–Macaulay is inherited by links. Thus, WA_n is homotopy Cohen–Macaulay as well. ∎

Proof of Proposition 4.23 We first check that ZA_n is bounded and graded. Since ZA_n is the dual of $\widehat{\mathrm{WA}}_n$, this is equivalent to showing that $\widehat{\mathrm{WA}}_n$ is bounded and graded. Boundedness is clear: $\widehat{\mathrm{WA}}_n$ has a global maximum by construction and a global minimum since WA_n does. The fact that $\widehat{\mathrm{WA}}_n$ is graded follows directly from Observation 4.7(g).

Let us now prove that ZA_n has a recursive atom ordering. Since WA_n is defined as a subposet of WO_{n+1}, we can view ZA_n as a subposet of ZO_{n+1}, with the same global minimum, call it $\widehat{0}$. Clearly, the subposet WA_n of WO_{n+1} is closed under taking lower bounds, since a folding cannot increase the degree of a labeled vertex, so the subposet ZA_n of ZO_{n+1} is closed under taking upper bounds. Hence, for any $T \preceq U$ in ZA_n, the interval $[T, U]$ is the same viewed in either ZA_n or ZO_{n+1}.

Denote by $\mathbf{0}$ the unique tree in WO_{n+1} of degree 0. Recall that it is the global minimum of WO_{n+1} and thus the global maximum of ZO_{n+1}. By [BMMM01, Lemmas 2.5 and 2.7], for every $\widehat{0} \neq T \in \mathrm{ZO}_{n+1}$, any ordering of the atoms of $[T, \mathbf{0}]$ in ZO_{n+1} is a recursive atom ordering, so by the previous paragraph, the same holds in ZA_n. In other words, the first condition in the definition of recursive atom ordering will always be automatically satisfied.

As for the second condition, by Observation 4.7(f)(g), the posets ZA_n and ZO_{n+1} are both graded of the same length (namely $n-1$) and the atoms of ZA_n are precisely the atoms of ZO_{n+1} that lie in ZA_n. Now we claim that the recursive atom ordering for ZO_{n+1} given in [BMMM01] induces an ordering of the atoms of ZA_n that is also a recursive atom ordering. Say the atoms of ZA_n are ordered $T_1, \dots, T_\ell$. Let $i < j$ and say $T_i, T_j \preceq U$ for some $U \in \mathrm{ZA}_n$. Working in ZO_{n+1}, where we already know we have a recursive atom ordering, we get that there exist $i' < j$ and $V \preceq U$ such that V covers $T_{i'}$ and T_j. Since ZA_n is closed under upper bounds and $T_j \preceq V$, we know that $V \in \mathrm{ZA}_n$. The last thing to do is show that $T_{i'} \in \mathrm{ZA}_n$.

The proof of [BMMM01, Theorem 5.12] essentially provides an algorithm for constructing $T_{i'}$ from V. In that proof, the elements of ZO_{n+1} are viewed as trees rooted at 1 and drawn upside down. According to this graphical representation, foldings are naturally divided into two types, called lifts and merges, and the corresponding unfoldings are called drops and splits (see [BMMM01, Definition 5.9], and also Figure 2 for examples). The proof shows that $T_{i'}$ can be obtained from V by a split (not a drop).

On the other hand, if we view V as an unrooted tree, then any unfolding at a leaf not labeled 1 is a drop. Since $V \in \mathrm{ZA}_n$, the vertex $n+1$ in V is a leaf not labeled 1. Hence, the unfolding that produces $T_{i'}$ from V is not based at $n+1$, so $n+1$ remains a leaf in $T_{i'}$ and hence $T_{i'} \in \mathrm{ZA}_n$. ∎

We are now ready to prove our main result.

Proof of parts (a) and (c) of Theorem 1.1 The proofs of these two parts are analogous, so we will only prove (a), asserting that $\Sigma^{n-2}(\mathrm{PSAut}_n)$ is dense in $\mathbb{S}(\mathrm{PSAut}_n)$. Let χ be any character of PSAut_n such that $\chi|_H \neq 0$ for all nontrivial $H \in \mathcal{HA}_n$. Let us verify the hypotheses of Theorem 4.3, to confirm that $[\chi] \in \Sigma^{n-2}(\mathrm{PSAut}_n)$. We

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

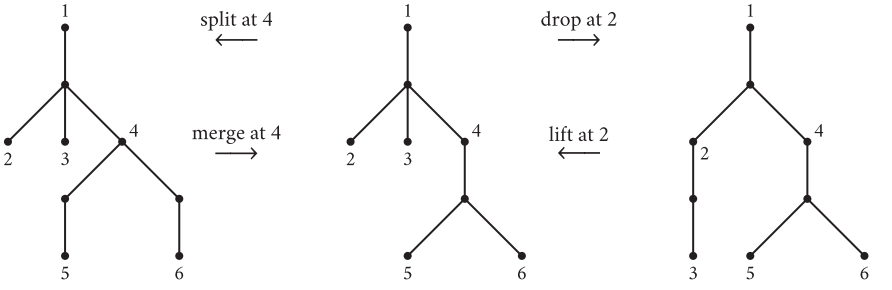


Figure 2: An example of a split/merge, and a drop/lift.

know that PSAut_n and all the $H_A(T)$ are of type F_∞, hence F_{n-2}, and that $\mathcal{HA}_n$ is non-empty, finite, intersection-closed (Proposition 4.14), and ∞-generating, hence $(n-2)$-generating (Proposition 4.17). Since $\chi|_H \neq 0$ for all nontrivial $H \in \mathcal{HA}_n$, we have that $[\chi|_H] \in \Sigma^{n-2}(H)$ for all such H, since H is abelian. Finally, $\mathcal{F}_0(\mathcal{HA}_n)$ is $(n-3)$-connected by Proposition 4.19. All the hypotheses of Theorem 4.3 have been confirmed, and so we conclude that $[\chi] \in \Sigma^{n-2}(\mathrm{PSAut}_n)$.

The last thing to show is that the set of all such $[\chi]$ is dense. It is clear from the McCool presentation that for any nontrivial $H_A(T)$, the image of $H_A(T)$ in the real vector space $\mathrm{PSAut}_n/[\mathrm{PSAut}_n, \mathrm{PSAut}_n] \otimes \mathbb{R}$ has positive dimension. Thus, taking the dual, the subspace of $\mathrm{Hom}(\mathrm{PSAut}_n, \mathbb{R})$ consisting of characters that send $H_A(T)$ to 0 has positive codimension. A finite union of subspaces with positive codimension has a dense complement, so we are done. ∎

5 More on Σ^2 for McCool groups

Let us say that a character χ of PSAut_n (resp. PSOut_n) is *generic* if $\chi|_H \neq 0$ for all nontrivial $H \in \mathcal{HA}_n$ (resp. $H \in \mathcal{HO}_n$). It is not difficult to check that a character χ of PSAut_n is generic if and only if $\chi(\alpha_{I,j}) \neq 0$ for all (nontrivial) Whitehead generators $\alpha_{I,j}$ and likewise a character χ of PSOut_n is generic if and only if $\chi(\overline{\alpha_{I,j}}) \neq 0$ whenever $\overline{\alpha_{I,j}} \neq 1$ (recall that $\overline{\alpha_{I,j}}$ is the image of $\alpha_{I,j}$ in PSOut_n). The proof of Theorem 1.1 shows that $\Sigma^m(\mathrm{PSAut}_n)$ and $\Sigma^m(\mathrm{PSOut}_n)$ contain (the equivalence classes of) all generic characters for $m \leq n-2$ and $m \leq n-3$, respectively, but does not say anything about non-generic characters.

The precise computation of $\Sigma^1(\mathrm{PSAut}_n)$ in [OK00] (see Theorem 5.1 below) suggests that $\Sigma^m(\mathrm{PSAut}_n)$ should be much larger than the set of generic characters. In this section, we will pose some specific questions about the structure of $\Sigma^m(\mathrm{PSAut}_n)$ and then answer one of these questions for $m = 2$. For brevity of exposition, we will limit our discussion to PSAut_n, but all the results and questions in this section have natural counterparts for PSOut_n.

Theorem 5.1 (Orlandi-Korner) *Let $n \geq 2$ and let χ be a nonzero character of* PSAut_n. *Then $[\chi] \notin \Sigma^1(\mathrm{PSAut}_n)$ if and only if one of the following holds:*

(1) *There exist distinct indices i, j such that $\chi(\alpha_{\ell,m}) = 0$ whenever $\{\ell, m\} \neq \{i, j\}$ as sets.*

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

(2) *There exist distinct indices* i, j, k *such that* $\chi(\alpha_{\ell,m}) = 0$ *whenever* $\{\ell, m\} \not\subseteq \{i, j, k\}$ *and in addition* $\chi(\alpha_{i,k}) + \chi(\alpha_{j,k}) = 0$, $\chi(\alpha_{i,j}) + \chi(\alpha_{k,j}) = 0$, *and* $\chi(\alpha_{j,i}) + \chi(\alpha_{k,i}) = 0$.

Define the *support* $\mathrm{supp}(\chi)$ of a character χ of PSAut_n to be the set of all (i, j) such that $\chi(\alpha_{i,j}) \neq 0$. Thus, Theorem 5.1 implies in particular that the classes of characters whose support has at least four elements always lie in $\Sigma^1(\mathrm{PSAut}_n)$, and one may ask if a similar phenomenon holds for other BNSR invariants.

Question 5.2 *Assume that* $m \leq n-2$ *and let* χ *a nonzero character of* PSAut_n.

(a) *Does* $\Sigma^m(\mathrm{PSAut}_n)$ *always contain* $[\chi]$ *whenever* $|\mathrm{supp}(\chi)|$ *is largest possible, that is,* $\chi(\alpha_{i,j}) \neq 0$ *for all* $i \neq j$*?*
(b) *Does there exist a constant* $C(m)$ *and a subset* $S(m,n)$ *of* $\{(i,j) \mid 1 \leq i \neq j \leq n\}$ *with* $|S(m,n)| = C(m)$ *such that* $\Sigma^m(\mathrm{PSAut}_n)$ *always contains* $[\chi]$ *whenever* $\mathrm{supp}(\chi)$ *contains* $S(m,n)$*?*
(c) *Does there exist a constant* $C(m)$ *such that* $\Sigma^m(\mathrm{PSAut}_n)$ *always contains* $[\chi]$ *whenever* $|\mathrm{supp}(\chi)| > C(m)$*?*

Clearly, a positive answer to any part of Question 5.2 yields a positive answer to the preceding parts. We also note that the conclusion of Theorem 1.1 as well as each of the conditions (a)–(c) above yields a simple geometric condition on the set $\mathbb{S}(\mathrm{PSAut}_n) \backslash \Sigma^m(\mathrm{PSAut}_n)$, the complement of $\Sigma^m(\mathrm{PSAut}_n)$ in the character sphere of PSAut_n.

We first recall that $\mathbb{S}(\mathrm{PSAut}_n)$ has dimension $n(n-1)-1$ and fix $m \leq n-2$. Theorem 1.1 implies that $\mathbb{S}(\mathrm{PSAut}_n) \backslash \Sigma^m(\mathrm{PSAut}_n)$ is contained in the union of finitely many spheres of codimension 1, but the number of spheres grows exponentially with n. A positive answer to (a) would reduce the number of spheres in the union to $n(n-1)$. Since $\Sigma^m(\mathrm{PSAut}_n)$ is invariant under a natural action of the symmetric group S_n, it is easy to show that (b) would imply that $\mathbb{S}(\mathrm{PSAut}_n) \backslash \Sigma^m(\mathrm{PSAut}_n)$ is contained in the union of finitely many spheres of dimension $2C(m)n$. Finally, (c) would imply that the dimensions of the spheres can be bounded by a function of m (independent of n).

The main goal of this section is to answer Question 5.2(b) in the affirmative for $m = 2$ and $n \geq 10$. We will show that the set $S(2,n) = \{(1,2),(3,4),(5,6),(7,8),(9,10)\}$ has the required property; in other words, we will prove the following.

Proposition 5.3 *Assume that* $n \geq 10$. *Then* $[\chi] \in \Sigma^2(\mathrm{PSAut}_n)$ *whenever* $\chi(\alpha_{i,j}) \neq 0$ *for all* $(i,j) \in \{(1,2),(3,4),(5,6),(7,8),(9,10)\}$.

We will prove Proposition 5.3 using the following general criterion.

Theorem 5.4 (Commutation Σ^2-criterion for quotients) *Let* $Q = \langle X \mid R \rangle$ *be a finitely presented group, let* R_1 *be a subset of* R, *and* $G = \langle X \mid R_1 \rangle$ *(so that there is a natural surjection from* G *to* Q*). Let* χ *be a character of* Q *such that:*

(a) $[\chi] \in \Sigma^2(G)$ *(we denote the character of* G *induced by* χ *by the same symbol), and*
(b) *for every* $r \in R \backslash R_1$, *there exists* $g_r \in G$ *with* $\chi(g_r) \neq 0$ *such that* g_r *commutes with the image of* r *in* G.

Then $[\chi] \in \Sigma^2(Q)$.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

Theorem 5.4 is a straightforward consequence of a criterion of Meinert (see Theorem A.1 in Appendix A) established in his Ph.D. thesis [Mei93] (see also a remark after the proof of Corollary 2.8 in [Mei97], where Theorem A.1 is stated in a different language). Since [Mei93] is not easily accessible, we will provide a proof of Theorem A.1 in Appendix A, where we will also explain why it implies Theorem 5.4.

Remark 5.5 Theorem 5.4 is particularly convenient to apply to a group $Q = \langle X \mid R \rangle$ in the case when the majority of relations in R are of the form $[x, y] = 1$ for $x, y \in X$; call such relations *RAAG-like*. In this case, we can let R_1 be a set of some or all of the RAAG-like relations from R, so that $G = \langle X \mid R_1 \rangle$ is a RAAG, and therefore $\Sigma^2(G)$ is completely determined by the main result of [MMV98] (see Theorem 5.6 below). This allows us to get a handle on condition (a). As for condition (b), a natural way to ensure this holds is to require that for each relator $r \in R \backslash R_1$, there exists a generator $x_r \in X$ such that, for every generator x appearing in r, the (RAAG-like) relation $[x_r, x] = 1$ holds and lies in R_1. Indeed, in this case, if we set $g_r = x_r$, then condition (b) holds for any χ satisfying $\chi(x_r) \neq 0$. The more RAAG-like relations from R are included in R_1, the easier it is to satisfy (b); however (unlike for Σ^1), including too many of the RAAG-like relations in R_1 can affect whether (a) holds. Thus, in practice, it can be a balancing act to find the most useful R_1.

Theorem 5.6 below describes the $m = 1$ and $m = 2$ cases of the full computation of the BNSR-invariant Σ^m for an arbitrary RAAG, proved by Meier, Meinert, and Van Wyk in [MMV98], and also see [BG99] for an alternative proof due to Bux and Gonzalez. Before stating Theorem 5.6, we introduce some standard notation and terminology. Let Γ be a finite graph with no loops or multiple edges, and let $\widehat{\Gamma}$ be the flag complex associated with Γ (i.e., the simplices of $\widehat{\Gamma}$ are precisely the cliques in Γ). Let A_Γ be the RAAG corresponding to Γ, that is, A_Γ has generators $\{x_v \mid v \text{ is a vertex of } \Gamma\}$ and defining relations $[x_v, x_w] = 1$ whenever v and w are adjacent.

Let χ be a nonzero character of A_Γ.

- A vertex v of Γ is called *living* (with respect to χ) if $\chi(v) \neq 0$ and *dead* otherwise.
- The *living subcomplex of* χ is the full subcomplex of $\widehat{\Gamma}$ spanned by the living vertices.
- A simplex σ of $\widehat{\Gamma}$ is called *dead* if all its vertices are dead.
- The *living link* of a simplex σ in $\widehat{\Gamma}$ is the full subcomplex of the link of σ spanned by its living vertices.

Theorem 5.6 *[BG99, MMV98] Let Γ and A_Γ be as above and χ a nonzero character of A_Γ. The following hold:*

(1) $[\chi] \in \Sigma^1(A_\Gamma)$ *if and only if the living subcomplex of χ is connected and every dead vertex is adjacent to a living vertex;*
(2) $[\chi] \in \Sigma^2(A_\Gamma)$ *if and only if*
   - *the living subcomplex of χ is simply-connected,*
   - *the living link of every dead vertex is connected and non-empty, and*
   - *the living link of every dead edge is non-empty.*

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

Condition (2) is often difficult to check in specific examples; however, its verification can sometimes be avoided entirely as there is a natural condition on Γ which guarantees that $\Sigma^1(A_\Gamma) = \Sigma^2(A_\Gamma)$, namely, the following.

Definition 5.7 A graph Γ is called *chordal* if any cycle C of length ≥ 4 in Γ has a chord, that is, an edge of Γ connecting non-consecutive vertices of C.

Whenever Γ is chordal we have $\Sigma^1(A_\Gamma) = \Sigma^2(A_\Gamma)$ [MMV98, Corollary 6.1], which makes understanding $\Sigma^2(A_\Gamma)$ much easier in this case. We now use RAAGs corresponding to chordal graphs to prove our main result of this section, Proposition 5.3.

Proof of Proposition 5.3 Let $S = \{(1,2),(3,4),(5,6),(7,8),(9,10)\}$ and define the graph Γ as follows. The vertices of Γ are all the pairs (i,j) with $1 \leq i \neq j \leq n$. Two distinct vertices (i,j) and (k,ℓ) are adjacent if and only if at least one of them lies in S and $[\alpha_{i,j}, \alpha_{k,l}] = 1$ in PSAut_n. Note that

(i) any two vertices in S are adjacent to each other,
(ii) every vertex of Γ is adjacent to a vertex in S, and
(iii) any vertex of Γ not in S can only be adjacent to vertices that are in S.

If C is a cycle in Γ, then either all vertices of C are in S or there is a vertex of C outside of S, in which case both of its neighbors must be in S by property (iii); either way, by (i), C contains a chord as soon as its length is at least 4. Thus, Γ is chordal.

Now let $Q = \mathrm{PSAut}_n$ (with the McCool presentation) and $G = A_\Gamma$. By construction, Q is a quotient of G. Let χ be a character of Q such that $\chi(\alpha_{i,j}) \neq 0$ whenever $(i,j) \in S$. We need to check that the induced character of G (also denoted by χ) satisfies the hypotheses of Theorem 5.4. The living subcomplex of χ contains S and hence is connected by properties (i) and (ii) above. Likewise, property (ii) implies that every dead vertex is adjacent to a living vertex. Hence, by Theorem 5.6(1), $[\chi]$ lies in $\Sigma^1(G)$, and $\Sigma^1(G) = \Sigma^2(G)$ since Γ is chordal. Thus, we have confirmed condition (a) of Theorem 5.4. We now check (b); actually, we will show that it holds for all $r \in R$, where R is the set of defining relators in the McCool presentation. This is because any $r \in R$ involves at most four indices, but S has five elements, so there exists $(k,\ell) \in S$ such that neither k nor ℓ appears in the subscript of any generator used in r, whence the corresponding generator $x_{(k,\ell)}$ of G commutes (in G) with every generator used in r. ∎

A A proof of Meinert's Σ^2-criterion

The goal of this appendix is to state and prove a criterion of Meinert established in his Ph.D. thesis [Mei93, Satz 5.2.1] and that is the key to proving Theorem 5.4.

In order to state Meinert's criterion, we need the notion of a generating set for groups with operators. Let H be a group, and suppose we are given a set Φ together with a map $\Phi \to \mathrm{Aut}(H)$. We will say that H is *finitely generated as a Φ-group* if there exists a finite subset S of H such that the smallest Φ-invariant subgroup of H containing S is H itself. If H is a normal subgroup of some group G and Φ is a subset of G, we can view H as a Φ-group where Φ acts on H by conjugation.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

Theorem A.1 (Meinert's Σ^2-criterion for quotients) *Let $1 \to N \to G \to Q \to 1$ be a short exact sequence of groups, with G and Q finitely presented. Let χ be a nonzero character of Q. Assume that*

(i) $[\chi] \in \Sigma^2(G)$ *(as before we denote the character of G induced by χ by the same symbol);*

(ii) *N is finitely generated as a $G_{\chi \geq 0}$-group, where $G_{\chi \geq 0} = \{g \in G \mid \chi(g) \geq 0\}$.*

Then $[\chi] \in \Sigma^2(Q)$.

This criterion should be compared to the classical fact that $[\chi] \in \Sigma^1(G)$ always implies $[\chi] \in \Sigma^1(Q)$ [BNS87, Proposition 3.3].

Let us first deduce Theorem 5.4 from Theorem A.1.

Proof of Theorem 5.4 Assume that G and Q are as in Theorem 5.4. Then we have the short exact sequence $1 \to N \to G \to Q \to 1$, where N is the normal closure of the image of $R \backslash R_1$ in G. Condition (i) in Theorem A.1 holds by assumption, so we just need to check (ii). By definition, any element $n \in N$ can be written as $n = \prod_{i=1}^{k} a_i r_i^{\pm 1} a_i^{-1}$, where $r_i \in R \backslash R_1$ and $a_i \in G$ (here we are equivocating between r_i and its image in G). Recall that by the hypotheses of Theorem 5.4, for each i, there exists g_i such that $\chi(g_i) \neq 0$ and $[g_i, r_i] = 1$ in G. The latter condition means that in the above expression for n, we can replace a_i by $a_i g_i^{m_i}$ for any integer m_i, and since $\chi(g_i) \neq 0$, after doing so we can assume that $a_i \in G_{\chi \geq 0}$. But this means that N is generated as a $G_{\chi \geq 0}$-group by the finite set $R \backslash R_1$, so hypothesis (ii) in Theorem A.1 holds, as desired. ∎

We now make some preparations for the proof of Theorem A.1. We will start with some terminology and notation. Let X be a set. As usual, by $F(X)$, we will denote the free group with basis X. By $\Omega(X)$, we will denote the set of all finite words in the alphabet $X \sqcup X^{-1}$ (not necessarily reduced). We can think of $F(X)$ as both a subset and a quotient of $\Omega(X)$ (the quotient map $\Omega(X) \to F(X)$ is a homomorphism of monoids). While no non-empty word in $\Omega(X)$ is invertible in $\Omega(X)$, there is a natural involution $w \mapsto w^{-1}$ on $\Omega(X)$: if $w = x_1 \dots x_n \in \Omega(X)$, with $x_i \in X^{\pm 1}$ for all i, we set $w^{-1} = x_n^{-1} \dots x_1^{-1}$, where by convention $(x^{-1})^{-1} = x$ for all $x \in X$.

Let R be a subset of $\Omega(X)$. By $R^{F(X)}$, we will denote the normal closure of the image of R in $F(X)$. By $R^{\Omega(X)}$, we will denote the set of all words of the form $\prod_{i=1}^{k} a_i r_i a_i^{-1}$, where $a_i \in \Omega(X)$ and $r_i \in R^{\pm 1}$. Thus, $R^{\Omega(X)}$ maps onto $R^{F(X)}$ under the natural projection $\Omega(X) \to F(X)$.

Definition A.2 Given a word $w = x_1 \dots x_n \in \Omega(X)$ with $x_i \in X^{\pm 1}$ for all i, by a *prefix* of w, we will mean a subword of the form $x_1 \dots x_k$ with $k \leq n$. We do allow $k = 0$.

If X is a generating set of some group G, any character of G naturally induces a character of the group $F(X)$ and a character of the monoid $\Omega(X)$ (for simplicity, we will not introduce separate notation for those induced characters).

Given $w \in \Omega(X)$ and a character χ of $\Omega(X)$, define

$$\chi_{min}(w) = \min\{\chi(v) \mid v \text{ is a prefix of } w\}.$$

We will say that w is *χ-nonnegative* if $\chi_{min}(w) = 0$ (since we count the empty word as a prefix, $\chi_{min}(w)$ is never positive).

The following observation collects some easy properties of the χ_{min} function.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

Observation A.3 *Let X be a set, and let $u, v \in \Omega(X)$. The following hold:*

(a) $\chi_{min}(uv) = \min\{\chi_{min}(u), \chi(u) + \chi_{min}(v)\} \geq \chi_{min}(u) + \chi_{min}(v)$.
(b) *(special case of (a)). If $\chi(u) = 0$, then $\chi_{min}(uv) = \min\{\chi_{min}(u), \chi_{min}(v)\}$.*
(c) *If $\chi(u) = 0$, then $\chi_{min}(u^{-1}) = \chi_{min}(u)$.*

We can rephrase the usual, topological definition of $\Sigma^2(G)$ combinatorially in terms of χ_{min} as follows.

Lemma A.4 *Let $\langle X \mid R \rangle$ be a finite presentation of a group G, and let χ be a nonzero character of G. Then $[\chi] \in \Sigma^2(G)$ if and only if $[\chi] \in \Sigma^1(G)$ and*

(Σ^2_{comb}) *: there is a constant C such that if $w \in \Omega(X)$ is any χ-nonnegative relator of G, then there exists $w' \in R^{\Omega(X)}$ such that $w' = w$ in $F(X)$ and $\chi_{min}(w') \geq C$.*

(Note that we must have $C \leq 0$ since $\chi_{min}(w')$ is never positive.)

Proof Let K be the Cayley complex of G corresponding to the finite presentation $\langle X \mid R \rangle$. Let $h_\chi \colon K \to \mathbb{R}$ be as in the definition of Σ-invariants.

First recall that $\Sigma^2(G)$ is contained in $\Sigma^1(G)$, and for a character χ, we have $[\chi] \in \Sigma^1(G)$ if and only if $K^{h_\chi \geq t}$ is connected for some (or, equivalently, all) $t \in \mathbb{R}$. Thus, Lemma A.4 only deals with characters χ such that $[\chi] \in \Sigma^1(G)$, and in view of Remark 2.2(b), it suffices to show that the following are equivalent for a given constant $C \leq 0$:

(a) the inclusion $K^{h_\chi \geq 0} \to K^{h_\chi \geq C}$ induces the trivial map $\pi_1(K^{h_\chi \geq 0}, 1_G) \to \pi_1(K^{h_\chi \geq C}, 1_G)$ (note that 1_G lies in both $K^{h_\chi \geq 0}$ and $K^{h_\chi \geq 0}$);
(b) (Σ^2_{comb}) holds for C.

We will use the following standard fact about the fundamental groups of CW-complexes. Let Y be a path-connected CW-complex and fix a vertex v_0 of Y. Then every element of $\pi_1(Y, v_0)$ is represented by an edge path, and if p is any closed edge path starting from v_0, then p represents the identity in $\pi_1(Y, v_0)$ if and only if there exists another closed edge path p' starting from v_0 such that:

(1) p' and p represent the same element in $\pi_1(Y^1, v_0)$, where Y^1 is the 1-skeleton of Y;
(2) p' can be written as a product $\prod_{i=1}^{k} \alpha_i \beta_i \alpha_i^{-1}$, where for each i, the path α_i is an edge path from v_0 to some vertex v_i of Y and there exists a 2-cell R_i containing v_i such that β_i is a closed path which starts at v_i and traverses the boundary of R_i.

We are now ready to prove the equivalence of conditions (a) and (b) above. Below we will establish the implication "(a)⇒ (b)." The other implication can be proved by just reversing the argument.

Suppose (a) holds. Let $w \in \Omega(X)$ be a χ-nonnegative relator of G, and let p be the unique edge path in $\mathrm{Cay}(G, X)$ starting at 1_G whose label is w. Note that p is closed since w is a relator of G.

The assumption that w is χ-nonnegative means precisely that p lies in $K^{h_\chi \geq 0}$, so by our choice of C, the path p represents the identity in $\pi_1(K^{h_\chi \geq C}, 1_G)$. Thus, there exists a closed edge path p' in $Y = K^{h_\chi \geq C}$ satisfying (1) and (2) above. Let $w' \in \Omega(X)$ denote the label of p'. Condition (1) means precisely that $w = w'$ in $F(X)$, and condition (2)

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

means precisely that $w' \in R^{\Omega(X)}$ with $\chi_{min}(w') \geq C$. Thus, w' satisfies the conclusion of (Σ^2_{comb}), and therefore (b) holds. ∎

We can now prove Meinert's criterion.

Proof of Theorem A.1 To reiterate the setup, we have a short exact sequence of groups $1 \to N \to G \to Q \to 1$ with G and Q finitely presented, and a nonzero character $\chi\colon Q \to \mathbb{R}$, with χ also denoting the induced character of G. We are assuming that $[\chi] \in \Sigma^2(G)$ and that N is finitely generated as a $G_{\chi\geq 0}$-group. Our goal is to prove that $[\chi] \in \Sigma^2(Q)$.

Since $[\chi] \in \Sigma^1(G)$, we clearly have $[\chi] \in \Sigma^1(Q)$ [BNS87, Proposition 3.3]. Hence, by Lemma A.4, it suffices to show that χ considered as a character of Q satisfies (Σ^2_{comb}).

Fix a finite presentation $\langle X \mid R\rangle$ of G. Choose a finite subset R_Q of $\Omega(X)$ whose image in G generates N as a $G_{\chi\geq 0}$-group, and set

$$C_Q = \min\{\chi_{min}(r) \mid r \in R_Q\} \quad \text{and} \quad C_X = \max\{\chi(x) \mid x \in X \sqcup X^{-1}\}.$$

Note that $C_Q \leq 0$ and $C_X > 0$. Next, let $C_G \leq 0$ be a constant for χ considered as a character of G in (Σ^2_{comb}), and set

$$C = C_Q - C_X + C_G.$$

Note that $C < C_Q$.

We will show that (Σ^2_{comb}) holds for χ as a character of Q using this constant C and the finite presentation $Q = \langle X \mid R \cup R_Q\rangle$. Let $w \in \Omega(X)$ be a χ-nonnegative relator of Q. The image of w in G lies in N, so we can choose $v \in \Omega(X)$ with $v = w$ in G such that $v = \prod_{i=1}^{k} b_i r_i b_i^{-1}$, for some $r_i \in R_Q^{\pm 1}$ and b_i satisfying $\chi(b_i) \geq 0$ for each i. Since $[\chi] \in \Sigma^1(G)$, without loss of generality, $\chi_{min}(b_i) \geq 0$ for all i; by Observation A.3(b), this ensures that $\chi_{\min}(v) \geq C_Q$. Now set $u = wv^{-1}$, so u is a relator of G, and since $\chi_{min}(w) = 0$ by assumption, we compute using Observation A.3(b)(c) that

$$\chi_{min}(u) \geq \min\{\chi_{min}(w), \chi_{min}(v^{-1})\} = \chi_{min}(v^{-1}) = \chi_{min}(v) \geq C_Q.$$

Next choose $x \in X \sqcup X^{-1}$ with $\chi(x) > 0$, and choose $n \geq 0$ minimal such that $\chi(x^n) \geq -C_Q$. By minimality, $\chi(x^n) < -C_Q + C_X$. Set $u' = x^n u x^{-n}$, so u' is a χ-nonnegative relator of G. Since $[\chi] \in \Sigma^2(G)$, we can choose $u'' \in R^{\Omega(X)}$ such that $u' = u''$ in $F(X)$ and $\chi_{min}(u'') \geq C_G$. Say $u'' = \prod_{j=1}^{l} a_j r_j a_j^{-1}$ for some $a_j \in \Omega(X)$ and $r_j \in R^{\pm 1}$. Since $\chi(a_j r_j a_j^{-1}) = 0$ for all j, by Observation A.3(b), we get

$$\min_j\{\chi_{min}(a_j r_j a_j^{-1})\} = \chi_{min}(u'') \geq C_G.$$

But now we can rewrite w as an element of $F(X)$ as follows:

$$w = uv = (x^{-n}u'x^n)v = (x^{-n}u''x^n)v = \left(\prod_{j=1}^{l} x^{-n} a_j r_j a_j^{-1} x^n\right) \cdot v.$$

Let $w' \in \Omega(X)$ denote this last word. Since $v \in R_Q^{\Omega(X)}$, we have that w' lies in $(R \cup R_Q)^{\Omega(X)}$ and $w = w'$ in $F(X)$. The only thing remaining is to prove that $\chi_{min}(w') \geq C$.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

Using Observation A.3(b) again, we have

$$\chi_{min}(w') = \min\{\chi_{min}(v), \chi_{min}(x^{-n}a_j r_j a_j^{-1} x^n) \mid 1 \leq j \leq l\}.$$

If $\chi_{min}(w') = \chi_{min}(v)$, then $\chi_{min}(w') \geq C_Q > C$ and we are done. Now suppose that $\chi_{min}(w') = \chi_{min}(x^{-n}a_j r_j a_j^{-1} x^n)$. By Observation A.3(a), we have

$$\chi_{min}(x^{-n}a_j r_j a_j^{-1} x^n) \geq \chi_{min}(x^{-n}) + \chi_{min}(a_j r_j a_j^{-1}) + \chi_{min}(x^n).$$

By construction, we have $\chi_{min}(x^{-n}) > C_Q - C_X$, $\chi_{min}(x^n) = 0$, and $\chi_{min}(a_j r_j a_j^{-1}) \geq C_G$, so putting everything together, we get $\chi_{min}(w') \geq C_Q - C_X + C_G = C$, as desired. ∎

Acknowledgements The authors are grateful to Sam Hughes, Dawid Kielak, and Lorenzo Ruffoni for helpful discussions and pointers to references. The authors also thank the referee for many helpful suggestions.

References

[AH93] H. Abels and S. Holz, *Higher generation by subgroups*. J. Algebra **160**(1993), no. 2, 310–341.

[AGMP25] P. Ardaiz-Galé and C. Martínez-Pérez, *The McCullough–Miller complex for right angled Artin groups*. Preprint. arXiv:2506.03377.

[BWC07] J. C. Baez, D. K. Wise, and A. S. Crans, *Exotic statistics for strings in 4D BF theory*. Adv. Theor. Math. Phys. **11**(2007), no. 5, 707–749.

[BT68] G. Baumslag and T. Taylor, *The centre of groups with one defining relator*. Math. Ann. **175**(1968), 315–319.

[BG03] R. Bieri and R. Geoghegan, *Connectivity properties of group actions on non-positively curved spaces*. Mem. Amer. Math. Soc. **161**(2003), no. 765, xiv+83.

[BGK10] R. Bieri, R. Geoghegan, and D. H. Kochloukova, *The sigma invariants of Thompson's group F*. Groups Geom. Dyn. **4**(2010), no. 2, 263–273.

[BNS87] R. Bieri, W. D. Neumann, and R. Strebel, *A geometric invariant of discrete groups*. Invent. Math. **90**(1987), no. 3, 451–477.

[BR88] R. Bieri and B. Renz, *Valuations on free resolutions and higher geometric invariants of groups*. Comment. Math. Helv. **63**(1988), no. 3, 464–497.

[BW83] A. Björner and M. Wachs, *On lexicographically shellable posets*. Trans. Amer. Math. Soc. **277**(1983), no. 1, 323–341.

[BGCAMP22] R. Blasco-García, J. I. Cogolludo-Agustín, and C. Martínez-Pérez, *On the Sigma-invariants of even Artin groups of FC-type*. J. Pure Appl. Algebra **226**(2022), no. 7, Article no. 106984, 20 pp.

[BMMM01] N. Brady, J. McCammond, J. Meier, and A. Miller, *The pure symmetric automorphisms of a free group form a duality group*. J. Algebra **246**(2001), no. 2, 881–896.

[BH13] T. E. Brendle and A. Hatcher, *Configuration spaces of rings and wickets*. Comment. Math. Helv. **88**(2013), no. 1, 131–162.

[Bro82] K. S. Brown, *Cohomology of groups*, volume 87 of Graduate Texts in Mathematics, Springer-Verlag, New York, NY, 1982.

[Bux04] K.-U. Bux, *Finiteness properties of soluble arithmetic groups over global function fields*. Geom. Topol. **8**(2004), 611–644.

[BFM+16] K.-U. Bux, M. G. Fluch, M. Marschler, S. Witzel, and M. C. B. Zaremsky, *The braided Thompson's groups are of type F_∞*. J. Reine Angew. Math. **718**(2016), 59–101. With an appendix by Zaremsky.

[BG99] K.-U. Bux and C. Gonzalez, *The Bestvina-Brady construction revisited: Geometric computation of Σ-invariants for right-angled Artin groups*. J. London Math. Soc. (2) **60**(1999), no. 3, 793–801.

[CGJ05] Y. Chen, H. H. Glover, and C. A. Jensen, *Proper actions of automorphism groups of free products of finite groups*. Int. J. Algebra Comput. **15**(2005), no. 2, 255–272.

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press

[Col89] D. J. Collins, *Cohomological dimension and symmetric automorphisms of a free group*. Comment. Math. Helv. **64**(1989), no. 1, 44–61.
[Dam17] C. Damiani, *A journey through loop braid groups*. Expo. Math. **35**(2017), no. 3, 252–285.
[EFMP24] M. Escartín-Ferrer and C. Martínez-Perez, *On the Σ-invariants of Artin groups satisfying the $K(\pi,1)$-conjecture*. J. Lond. Math. Soc. (2) **109**(2024), no. 1, Article no. e12861, 35 pp.
[FM12] B. Farb and D. Margalit, *A primer on mapping class groups*, volume 49 of Princeton Mathematical Series, Princeton University Press, Princeton, NJ, 2012.
[Gri13] J. T. Griffin, *Diagonal complexes and the integral homology of the automorphism group of a free product*. Proc. Lond. Math. Soc. (3) **106**(2013), no. 5, 1087–1120.
[HK24] S. Hughes and D. Kielak, *BNSR invariants and ℓ^2-homology*. Preprint. arXiv:2401.05545.
[LI24] C. L. Isenrich, *From the second BNSR invariant to Dehn functions of coabelian subgroups*. Preprint. arXiv:2404.12334.
[JMM06] C. Jensen, J. McCammond, and J. Meier, *The integral cohomology of the group of loops*. Geom. Topol. **10**(2006), 759–784.
[KMM15] N. Koban, J. McCammond, and J. Meier, *The BNS-invariant for the pure braid groups*. Groups Geom. Dyn. **9**(2015), no. 3, 665–682.
[KM22] D. H. Kochloukova and L. Mendonça, *On the Bieri-Neumann-Strebel-Renz Σ-invariants of the Bestvina-Brady groups*. Forum Math. **34**(2022), no. 3, 605–626.
[Lüc02] W. Lück, *L^2-invariants: Theory and applications to geometry and K-theory*, volume 44 of Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics], Springer-Verlag, Berlin, 2002.
[MM04] J. McCammond and J. Meier, *The hypertree poset and the l^2-Betti numbers of the motion group of the trivial link*. Math. Ann. **328**(2004), no. 4, 633–652.
[McC86] J. McCool, *On basis-conjugating automorphisms of free groups*. Canad. J. Math. **38**(1986), no. 6, 1525–1529.
[MM96] D. McCullough and A. Miller, *Symmetric automorphisms of free products*. Mem. Amer. Math. Soc. **122**(1996), no. 582, viii+97.
[MMV98] J. Meier, H. Meinert, and L. VanWyk, *Higher generation subgroup sets and the Σ-invariants of graph groups*. Comment. Math. Helv. **73**(1998), no. 1, 22–44.
[MMV01] J. Meier, H. Meinert, and L. VanWyk, *On the Σ-invariants of Artin groups*. Topology and its Applications **110**(2001), 71–81. Geometric topology and geometric group theory (Milwaukee, WI, 1997).
[Mei93] H. Meinert, *Die höheren geometrischen Invarianten Σ^m von Gruppen via Operationen auf CW-Komplexen und der Beweis der Σ^m-Vermutung für metabelsche Gruppen endlichen Prüfferr-ranges*. Ph.D. thesis, Universität Frankfurt am, 1993.
[Mei97] H. Meinert, *Actions on 2-complexes and the homotopical invariant Σ^2 of a group*. J. Pure Appl. Algebra **119**(1997), no. 3, 297–317.
[OK00] L. A. Orlandi-Korner, *The Bieri-Neumann-Strebel invariant for basis-conjugating automorphisms of free groups*. Proc. Amer. Math. Soc. **128**(2000), no. 5, 1257–1262.
[Ren88] B. Renz, *Geometrische Invarianten und Endlichkeitseigenschaften von Gruppen*. Ph.D. thesis, Universität Frankfurt, 1988.
[Str12] R. Strebel, *Notes on the Sigma invariants*. Preprint. arXiv:1204.0214.
[Zar17] M. C. B. Zaremsky, *Separation in the BNSR-invariants of the pure braid groups*. Publ. Mat. **61**(2017), no. 2, 337–362.
[Zar18] M. C. B. Zaremsky, *Symmetric automorphisms of free groups, BNSR-invariants, and finiteness properties*. Michigan Math. J. **67**(2018), no. 1, 133–158.

University of Virginia, USA
e-mail: ershov@virginia.edu

University at Albany (SUNY), USA
e-mail: mzaremsky@albany.edu

https://doi.org/10.4153/S0008414X26102089 Published online by Cambridge University Press