SEPARATION IN THE BNSR-INVARIANTS OF THE PURE BRAID GROUPS

Matthew C. B. Zaremsky

Abstract: We inspect the BNSR-invariants $\Sigma^m(P_n)$ of the pure braid groups P_n, using Morse theory. The BNS-invariants $\Sigma^1(P_n)$ were previously computed by Koban, McCammond, and Meier. We prove that for any $3 \leq m \leq n$, the inclusion $\Sigma^{m-2}(P_n) \subseteq \Sigma^{m-3}(P_n)$ is proper, but $\Sigma^\infty(P_n) = \Sigma^{n-2}(P_n)$. We write down explicit character classes in each relevant $\Sigma^{m-3}(P_n) \setminus \Sigma^{m-2}(P_n)$. In particular we get examples of normal subgroups $N \leq P_n$ with $P_n/N \cong \mathbb{Z}$ such that N is of type F_{m-3} but not F_{m-2}, for all $3 \leq m \leq n$.

2010 Mathematics Subject Classification: Primary: 20F65; Secondary: 57M07, 20F36.

Key words: Braid group, BNSR-invariant, finiteness properties.

Introduction

The Σ-invariants of a group G are a sequence of geometric objects $\Sigma^m(G)$ $(m \in \mathbb{N})$, defined whenever G is of type F_m, that encode a great deal of information about G and its subgroups. The first invariant, Σ^1, was introduced by Bieri, Neumann, and Strebel [**BNS**] and is also called the BNS-invariant. The higher invariants, Σ^2, Σ^3, and so forth, culminating in Σ^∞, were subsequently introduced by Bieri and Renz [**BR**], and are also known as BNSR-invariants. Once one knows the BNSR-invariants of a group, one gets a complete classification of which of its coabelian subgroups have which finiteness properties. (A *coabelian* subgroup is a normal subgroup with abelian quotient.)

Once the finiteness properties of a group are known (e.g., whether it is finitely generated, finitely presented, type F_m, etc.), the group's BNSR-invariants are a very natural next question. However, in general the BNSR-invariants of a group are notoriously difficult to compute. A complete computation has been done only for very few families of "interesting" groups. The main example where the problem is interesting, difficult, and totally solved is the case of right-angled Artin groups, whose BNSR-invariants were computed by Meier, Meinert, and

VanWyk [**MMV1**] and, independently, by Bux and Gonzalez [**BG**]. There are also some results for general Artin groups [**MMV2**]. The Morse-theoretic approach we take here is related to the methods used in [**WZ, Zar**] to compute BNSR-invariants of generalizations of Thompson's groups.

The braid groups B_n and pure braid groups P_n are extremely well studied families of groups. They arise in a variety of contexts, for example in the study of Artin groups, mapping class groups and knot theory. The braid group B_n is of type F, meaning it admits a compact classifying space, so from the point of view of finiteness properties B_n is as "good" as it can possibly be. Since P_n has finite index in B_n, it also has this property. From the point of view of BNSR-invariants though, the pure braid groups are vastly more complicated than the braid groups. The abelianization of B_n, for any $n \geq 2$, is just $\mathbb{Z}$, so the $\Sigma^m(B_n)$ all live in a 0-sphere, and it is easy to show that the whole 0-sphere equals $\Sigma^\infty(B_n)$ (Corollary 3.4). In particular, every coabelian subgroup of B_n is of type F_∞. The abelianization of P_n, however, is $\mathbb{Z}^{\binom{n}{2}}$, so the $\Sigma^m(P_n)$ all lie in an $(\binom{n}{2} - 1)$-sphere. With such a huge sphere, figuring out which parts of it do or do not lie in which BNSR-invariants is massively more complicated. Correspondingly, it is much more difficult to figure out the finiteness properties of coabelian subgroups of P_n.

The first Σ-invariant, that is the BNS-invariant $\Sigma^1(P_n)$, was computed by Koban, McCammond, and Meier [**KMM**]. In particular for $n \geq 3$ there are parts of the character sphere that do not lie in $\Sigma^1(P_n)$, and hence there are coabelian subgroups of P_n (for example the commutator subgroup itself) that are not finitely generated, in contrast to the B_n case. Since $P_3 \cong F_2 \times \mathbb{Z}$, every finitely generated coabelian subgroup is already finitely presented and even of type F_∞. For $n > 3$ though, it was unclear whether there exist coabelian subgroups that are finitely generated but not finitely presentable.

In this paper we use Morse theory to find regions of the character sphere of P_n that lie in $\Sigma^m(P_n)$ for various m, and regions that do not. Our main results are:

Theorem 5.7. *For any* $3 \leq m \leq n$*, the inclusion* $\Sigma^{m-2}(P_n) \subseteq \Sigma^{m-3}(P_n)$ *is proper.*

Theorem 3.9. *For any* $n \in \mathbb{N}$*,* $\Sigma^\infty(P_n) = \Sigma^{n-2}(P_n)$*.*

Consequences of this include that there exist coabelian subgroups of P_n, for any $n \geq 4$, that are finitely generated but not finitely presentable. More generally there exist coabelian subgroups of P_n that are

of type F_{m-3} but not F_{m-2}, for any $3 \leq m \leq n$. We give explicit examples of these; see Corollary 5.8. However, a coabelian subgroup of P_n of type F_{n-2} is automatically of type F_∞.

Our approach reduces the problem of proving a character class is or is not in a given BNSR-invariant of P_n to a sufficient condition on a finite complex PW_n. Hence it seems likely that our approach could reveal still more about the $\Sigma^m(P_n)$ in the future. We mention that even the problem of fully computing $\Sigma^2(P_4)$ (which equals $\Sigma^\infty(P_4)$) remains open.

This paper is organized as follows. In Section 1 we recall the BNSR-invariants and discuss some general results, and in Section 2 we set up our Morse-theoretic approach. In Section 3 we discuss the (pure) braid groups and their characters, and in Section 4 a complex on which they act. Finally in Section 5 we use Morse theory on this complex to derive our results about the BNSR-invariants of the pure braid groups.

1. BNSR-invariants

To define the BNSR-invariants $\Sigma^m(G)$ of a group G, we first need some background. A connected CW-complex Z is called a *classifying space* for G, or a $K(G,1)$, if $\pi_1(Z) \cong G$ and $\pi_k(Z) = 0$ for $k \geq 2$. We say G is of *type* F_n if it admits a $K(G,1)$ with compact n-skeleton. If G is of type F_n for all n we say it is of *type* F_∞. If there exists a compact $K(G,1)$ we say G is of *type* F.

Now suppose G is a group of type F_n and let Z be a $K(G,1)$ with compact n-skeleton. The universal cover $\widetilde{Z^{(n)}}$ of the n-skeleton is $(n-1)$-connected, and G acts freely and cocompactly on it. In fact, there is a converse to this: if G acts freely (or just properly) and cocompactly on an $(n-1)$-connected CW-complex, then G is of type F_n. This brings us to the definition of the BNSR-invariants.

Definition 1.1 (BNSR-invariants). Let G act properly cocompactly on an $(n-1)$-connected CW-complex Y (so G is of type F_n). Let $\chi \colon G \to \mathbb{R}$ be a *character* of G, i.e. a homomorphism to $\mathbb{R}$. There exists a continuous map $h_\chi \colon Y \to \mathbb{R}$ such that $h_\chi(g \cdot y) = \chi(g) + h_\chi(y)$ for all $g \in G$ and $y \in Y$. For $t \in \mathbb{R}$ let $Y_{\chi \geq t}$ be the full subcomplex of Y supported on those vertices y with $h_\chi(y) \geq t$. For non-trivial χ, let $[\chi]$ be the equivalence class of χ under scaling by positive real numbers. For $m \leq n$, the mth

BNSR-invariant $\Sigma^m(G)$ is defined to be

$$\Sigma^m(G) := \{[\chi] \mid (Y_{\chi \geq t})_{t \in \mathbb{R}} \text{ is essentially } (m-1)\text{-connected}\}.$$

The definition of Σ^m is admittedly quite dense. Let us unpack it a bit. The characters χ are elements of $\mathrm{Hom}(G, \mathbb{R})$, which is a vector space $\mathbb{R}^d$ for some d. The equivalence classes $[\chi]$ lie in the so called sphere at infinity S^{d-1} of $\mathbb{R}^d$, or *character sphere* $S(G)$. A character χ induces a *character height function* $h_\chi \colon Y \to \mathbb{R}$. The height function h_χ can be used to divide Y into regions, for example the region whose vertices have height greater than or equal to t. By varying t, we get a nested family of these "half spaces". Now if, for example, each half space were itself $(m-1)$-connected, then the definition says $[\chi]$ would lie in $\Sigma^m(G)$. Heuristically, $\Sigma^m(G)$ is a catalog of which half spaces are how highly connected. However, this is not quite right, as the presence of the word "essentially" indicates. Instead, we do not require each half space to be $(m-1)$-connected itself, but only that for all t there exists $-\infty < s \leq t$ such that the inclusion $Y_{\chi \geq t} \to Y_{\chi \geq s}$ induces the trivial map in π_k for $k \leq m-1$. (If the domain is already $(m-1)$-connected, then of course this will be the case.) For example, maybe $Y_{\chi \geq t}$ is not connected, but if all its components get connected up in $Y_{\chi \geq t-1}$ then that is enough to say $[\chi]$ is in $\Sigma^1(G)$.

It is not obvious from the definition, but $\Sigma^m(G)$ is well defined up to the various non-canonical aspects, e.g., the space Y and the height function h_χ. See for example [**Bux**, Definition 8.1]. As a remark, the definition there used the filtration by sets $h_\chi^{-1}([t, \infty))_{t \in \mathbb{R}}$, but thanks to cocompactness this filtration is essentially $(m-1)$ connected if and only if our filtration $(Y_{\chi \geq t})_{t \in \mathbb{R}}$ is.

One important application of BNSR-invariants is the following:

Citation 1.2 ([**BGK**, Theorem 1.1]). *Let G be a group of type F_m. Let H be any coabelian subgroup of G. Then H is of type F_m if and only if for every non-trivial character χ of G such that $\chi(H) = 0$, we have $[\chi] \in \Sigma^m(G)$.*

For example, if $H = \ker(\chi)$ for χ a *discrete* character of G, i.e. one with image $\mathbb{Z}$, then H is of type F_m if and only if $[\pm\chi] \in \Sigma^m(G)$.

Other important properties of the $\Sigma^m(G)$ are that they are all open subsets of $S(G)$, and that they are invariant under the natural action of $\mathrm{Aut}(G)$ on $S(G)$. We collect here some additional general results about BNSR-invariants that will be useful in the following sections.

Citation 1.3 ([**Mei**, Corollary 2.8]). *Let $G \xrightarrow{\pi} Q$ be a split epimorphism of groups. Let $\chi\colon Q \to \mathbb{R}$ be a character. Let $\widetilde{\chi}\colon G \to \mathbb{R}$ be $\widetilde{\chi} = \chi \circ \pi$. If $[\widetilde{\chi}] \in \Sigma^m(G)$ then $[\chi] \in \Sigma^m(Q)$.*

Lemma 1.4. *Let $G \xrightarrow{\pi} Q$ be an epimorphism of groups. Suppose that $\ker(\pi)$ is of type F_m. Let $\chi\colon Q \to \mathbb{R}$ be a character. Let $\widetilde{\chi}\colon G \to \mathbb{R}$ be $\widetilde{\chi} = \chi \circ \pi$. If $[\psi] \in \Sigma^m(Q)$ for all non-trivial characters ψ of Q such that $\psi(\ker(\chi)) = 0$, then $[\widetilde{\chi}] \in \Sigma^m(G)$.*

Proof: Consider the restriction π' of π to $\ker(\widetilde{\chi})$, so $\pi'\colon \ker(\widetilde{\chi}) \to \ker(\chi)$. Note that π' is surjective and $\ker(\pi') = \ker(\pi)$, so we have a short exact sequence
$$1 \to \ker(\pi) \to \ker(\widetilde{\chi}) \to \ker(\chi) \to 1.$$
We are assuming $[\psi] \in \Sigma^m(Q)$ for all non-trivial characters ψ of Q such that $\psi(\ker(\chi)) = 0$, so Citation 1.2 says that $\ker(\chi)$ is of type F_m. We are also assuming $\ker(\pi)$ is of type F_m, so in fact $\ker(\widetilde{\chi})$ is of type F_m. Again by Citation 1.2, we conclude that $[\widetilde{\chi}] \in \Sigma^m(G)$. □

This next result was called the Σ^m-criterion in [**MMV2**]; the homological version is Theorem 4.1 in [**BR**], and a proof of this homotopical version follows by mimicking that proof.

Citation 1.5. *With the notation from Definition 1.1, $[\chi] \in \Sigma^m(G)$ if and only if there exists a continuous, cellular G-equivariant map $\varphi\colon Y^{(m)} \to Y^{(m)}$ satisfying $h_\chi(\varphi(y)) > h_\chi(y)$ for all $y \in Y^{(m)}$.*

Corollary 1.6 (Center survives). *Let G be a group of type F_m. For a character $\chi\colon G \to \mathbb{R}$, if $\chi(Z(G)) \neq 0$ then $[\chi] \in \Sigma^m(G)$.*

This is proved in [**MMV2**, Lemma 2.1], but the proof is so short we may as well recreate it here.

Proof: Let Y be an $(m-1)$-connected complex on which G acts properly cocompactly. Choose $z \in Z(G)$ such that $\chi(z) > 0$. Let $\varphi\colon Y^{(m)} \to Y^{(m)}$ be the map $y \mapsto z.y$. Since z is central, this is G-equivariant, so by Citation 1.5, $[\chi] \in \Sigma^m(G)$. □

2. Morse theory

When trying to compute the BNSR-invariants of a group, Bestvina–Brady Morse theory [**BB**] can be a useful tool. Provided one sets things up correctly, Morse theory can reduce a global statement, like "this filtration is essentially $(m-1)$-connected", to local statements, like "the 'ascending' part of any vertex link is $(m-1)$-connected". In practice, the links may be finite simplicial complexes, which makes the latter question easier, in theory.

We will focus now on *affine cell complexes*. An affine cell complex is the quotient of a disjoint union of euclidean polytopes by an equivalence relation mapping each polytope injectively into the complex, such that the images of the polytopes, called *cells*, meet in faces (see [**BH**, Definition I.7.37]). Each cell of an affine cell complex Y carries an affine structure. The *star* $\mathrm{st}_Y v$ of a vertex v in Y is the subcomplex of Y consisting of cells that are faces of cells containing v. The *link* $\mathrm{lk}_Y v$ of v is the set of directions out of v into $\mathrm{st}_Y v$. The link is a simplicial complex, whose simplices consist of such directions into a given cell.

In [**BB**], Bestvina and Brady defined a *Morse function* on an affine cell complex Y to be a map $Y \to \mathbb{R}$ that is affine on cells, takes discretely many values on the vertices, and is non-constant on edges. When using Morse theory to compute BNSR-invariants, one often needs some more subtlety, for example if one is dealing with non-discrete character height functions. Our definition of *Morse function* here is similar to the one in [**WZ**, **Zar**]:

Definition 2.1 (Morse function)**.** Let Y be an affine cell complex and let $(h, f) \colon Y \to \mathbb{R} \times \mathbb{R}$ be a map such that the restrictions of h and f to any cell are affine functions. Suppose that there exists $\varepsilon > 0$ such that for any pair of adjacent vertices v and w, either $|h(v) - h(w)| \geq \varepsilon$, or else $h(v) = h(w)$ and $f(v) \neq f(w)$. If the set $f(Y^{(0)})$, considered under the usual total ordering of $\mathbb{R}$, has the property that any subset has a maximal element (that is, if the total ordering on $f(Y^{(0)})$ is an inverse well ordering), then we call (h, f) an *ascending-type Morse function*. If the set $f(Y^{(0)})$ with the usual total ordering induced by $\mathbb{R}$ has the property that any subset has a minimal element (so it is a well ordering), then we call (h, f) a *descending-type Morse function*. By a *Morse function* we mean either an ascending-type or descending type Morse function.

This requirement that the total ordering on $f(Y^{(0)})$ be an (inverse) well ordering is different from the definition of Morse function in [**WZ**, **Zar**]. There we just required $f(Y^{(0)})$ to be finite, which of course implies the present definition. Here we need this more robust definition, since our f (the function $\dot{\omega}$ in Section 5) will take infinitely many values.

Remark 2.2 (Classical Morse function). We will occasionally also deal with functions that are Morse functions in the sense of Bestvina–Brady, which in our language is a Morse function (h, f) where f is constant, say 0, and $h(Y^{(0)})$ is a discrete subset of $\mathbb{R}$.

Given a Morse function (h, f) on an affine cell complex Y, one can define the important notion of *ascending link*. First note that, on any

cell c, (h, f) achieves its maximum and minimum values on unique vertices of c. This is because h and f are affine on c, and are non-constant on all faces with positive dimension. Here we always use the lexicographic ordering on $\mathbb{R} \times \mathbb{R}$ to compare (h, f)-values. Define the *ascending star* st$\uparrow v$ of a vertex v in Y to be the subcomplex of st$_Y v$ consisting of those cells on which (h, f) achieves its minimum at v. Similarly the *descending star* st$\downarrow v$ is the subcomplex of cells whose (h, f)-maximum is at v. Now define the *ascending link* lk$\uparrow v$ to be the subcomplex of lk$_Y v$ consisting of those directions out of v into st$\uparrow v$. The *descending link* lk$\downarrow v$ is defined analogously. When all the cells are simplices, lk$_Y v$ can be identified with the subcomplex of st$_Y v$ consisting of those simplices not containing v, with lk$\uparrow v$ and lk$\downarrow v$ similarly viewed as subcomplexes.

Lemma 2.3 (Morse Lemma). *Let Y be an affine cell complex and $(h, f) \colon Y \to \mathbb{R} \times \mathbb{R}$ an ascending-type Morse function. For $t \in \mathbb{R}$ let $Y_{h \geq t}$ be the subcomplex of Y supported on vertices v with $h(v) \geq t$. Let $s < t$ (allow the case $s = -\infty$). If for all vertices v with $s \leq h(v) < t$ the ascending link $\mathrm{lk}\uparrow v$ with respect to (h, f) is $(m - 1)$-connected, then the inclusion $Y_{h \geq t} \to Y_{h \geq s}$ induces an isomorphism in π_k for $k \leq m - 1$ and a surjection in π_m.*

Now suppose instead that (h, f) is a descending-type Morse function. Let $Y_{h \leq t}$ be the subcomplex supported on vertices v with $h(v) \leq t$. Let $s > t$ (allow $s = \infty$). If for all vertices v with $t < h(v) \leq s$ the descending link $\mathrm{lk}\downarrow v$ with respect to (h, f) is $(m - 1)$-connected, then the inclusion $Y_{h \leq t} \to Y_{h \leq s}$ induces an isomorphism in π_k for $k \leq m - 1$ and a surjection in π_m.

Proof: The "descending" version can be converted into the "ascending" version by replacing (h, f) with $(-h, -f)$, so we just need to prove the ascending version. The proof should be compared to the proofs in [**BB**] of Lemma 2.5 and Corollary 2.6.

If we can prove the result assuming $t - s \leq \varepsilon$, then we will be done, thanks to induction (and compactness of spheres, for the $s = -\infty$ case). Let $Y_{s \leq h < t}$ be the subcomplex supported on vertices v with $s \leq h(v) < t$. In particular, $Y_{h \geq s}$ is obtained from $Y_{h \geq t}$ by attaching, in some order, the vertices of $Y_{s \leq h < t}$ along their relative links. The order needs to be an inverse well order, so that the set of vertices not yet attached has a unique maximum, i.e. a well defined "next" vertex to attach. Let $\prec$ be any inverse well order on the vertex set of $Y_{s \leq h < t}$ satisfying the property that if $f(v) < f(w)$ then $v \prec w$. Since (h, f) is an ascending-type Morse function, the set of f-values on vertices is itself inversely well ordered, so such a $\prec$ does indeed exist. When we write $v \prec w$, we mean that w gets attached before v, so v can "see" w but not vice versa.

The goal is now to show that when we attach a vertex, we do so along an $(m-1)$-connected relative link, so this induces an isomorphism in π_k for $k \leq m-1$ and a surjection in π_m, by the Mayer–Vietoris and Seifert–van Kampen Theorems. Then by transfinite induction, which we can use since $\prec$ is an inverse well order, we will get that the inclusion $Y_{h \geq t} \to Y_{h \geq s}$ also induces these types of maps, and we will be done.

When we attach the vertex v, we do so along a relative link $\mathrm{lk}_{\mathrm{rel}}\, v$ equal to the full subcomplex of $\mathrm{lk}_Y\, v$ supported on those vertices w with either $h(w) \geq t$, or else $s \leq h(w) < t$ and $v \prec w$. We now claim that this is exactly the ascending $\mathrm{lk}{\uparrow}v$. First note that for any w adjacent to v, Definition 2.1 says either $|h(v) - h(w)| \geq \varepsilon$, or else $h(v) = h(w)$ and $f(v) \neq f(w)$. For w of the first form, since $t - s \leq \varepsilon$ we know w cannot lie in $Y_{s \leq h < t}$. Hence $h(w) > h(v)$ if and only if $h(w) \geq t$, so such a w lies in $\mathrm{lk}{\uparrow}v$ if and only if it lies in $\mathrm{lk}_{\mathrm{rel}}\, v$. Now suppose w is of the second form, so $h(v) = h(w)$ and $f(v) \neq f(w)$. In particular $s \leq h(w) < t$. Now w is in $\mathrm{lk}{\uparrow}v$ if and only if $f(v) < f(w)$, and $w \in \mathrm{lk}_{\mathrm{rel}}\, v$ if and only if $v \prec w$, so we need to show that for such w, $f(v) < f(w)$ if and only if $v \prec w$. This follows from our construction of $\prec$, since we know $f(v) \neq f(w)$.

Having shown $\mathrm{lk}_{\mathrm{rel}}\, v = \mathrm{lk}{\uparrow}v$, we now know that $Y_{h \geq s}$ is obtained from $Y_{h \geq t}$ by coning off the ascending links of vertices in $Y_{s \leq h < t}$. By assumption, these are all $(m-1)$-connected, and so we are done. $\square$

As a corollary to the proof, we have:

Corollary 2.4. *With the same setup as the (ascending version of the) Morse Lemma, if additionally for all vertices v with $s \leq h(v) < t$ we have $\widetilde{H}_{m+1}(\mathrm{lk}{\uparrow}v) = 0$, then the inclusion $Y_{h \geq t} \to Y_{h \geq s}$ induces an injection in $\widetilde{H}_{m+1}$.*

With the setup of the descending version, a similar result holds with all the signs reversed and "ascending" changed to "descending".

Proof: In the proof of the Morse Lemma we saw that $Y_{h \geq s}$ is obtained from $Y_{h \geq t}$ by coning off the ascending links of vertices in $Y_{s \leq h < t}$, so this is immediate from the Mayer–Vietoris sequence. (The descending version follows similarly.) $\square$

One main use of the Morse Lemma is the case $s = -\infty$:

Corollary 2.5. *Let Y be an $(m-1)$-connected affine cell complex with an ascending-type Morse function (h, f). Suppose there exists N such that, for every vertex v of Y with $h(v) < N$, $\mathrm{lk}{\uparrow}v$ is $(m-1)$-connected. Then the filtration $(Y_{h \geq t})_{t \in \mathbb{R}}$ is essentially $(m-1)$-connected.*

The descending version holds as well, with all the signs reversed.

Proof: By the Morse Lemma, for any $r \leq N$ the inclusion $Y_{h \geq r} \to Y = Y_{h \geq -\infty}$ induces an isomorphism in π_k for $k \leq m - 1$. Since Y is $(m - 1)$-connected, so is $Y_{h \geq r}$. Now for any t we just need to choose $s = \min\{N, t\}$ and we get that the inclusion $Y_{h \geq t} \to Y_{h \geq s}$ induces the trivial map in π_k for $k \leq m - 1$, simply because $Y_{h \geq s}$ is $(m-1)$-connected. (The descending version follows similarly.) □

It is less straightforward to use Morse theory to prove that a filtration is *not* essentially m-connected. However, given some additional assumptions we can say something.

Proposition 2.6. *Let Y be an affine cell complex and let $(h, f) \colon Y \to \mathbb{R} \times \mathbb{R}$ be an ascending-type Morse function. Suppose there exists $N \in \mathbb{R}$ such that for every vertex v with $h(v) < N$ the ascending link $\mathrm{lk} \!\uparrow\! v$ is $(m - 1)$-connected and satisfies $\widetilde{H}_{m+1}(\mathrm{lk} \!\uparrow\! v) = 0$. Assume moreover that for all $M \in \mathbb{R}$ there exists a vertex v with $h(v) < M$ such that $\widetilde{H}_m(\mathrm{lk} \!\uparrow\! v) \neq 0$. If $\widetilde{H}_{m+1}(Y) = 0$, then the filtration $(Y_{h \geq t})_{t \in \mathbb{R}}$ is not essentially m-connected.*

If (h, f) is instead descending-type, then a similar result holds with all the signs reversed and "ascending" changed to "descending".

Proof: Suppose to the contrary that $(Y_{h \geq t})_{t \in \mathbb{R}}$ is essentially m-connected. Say $t < N$, and choose $-\infty < s \leq t$ such that the inclusion $Y_{h \geq t} \to Y_{h \geq s}$ induces the trivial map in π_k for $k \leq m$. Also, since $t < N$, this inclusion induces a surjection in these π_k by the Morse Lemma, so in fact $Y_{h \geq s}$ itself is m-connected, as are all $Y_{h \geq r}$ for $r \leq s$ (for the same reason). Now choose v such that $h(v) < s$ and $\widetilde{H}_m(\mathrm{lk} \!\uparrow\! v) \neq 0$. Since $\widetilde{H}_m(Y_{h \geq r}) = 0$ for all $r \leq s$, Mayer–Vietoris and Corollary 2.4 say that $\widetilde{H}_{m+1}(Y_{h \geq q}) \neq 0$ for any $q \leq h(v)$. But this includes $q = -\infty$, and $\widetilde{H}_{m+1}(Y) \neq 0$ contradicts our assumptions. (As usual, the descending version follows similarly.) □

For example, if all the $\mathrm{lk} \!\uparrow\! v$ are homotopy equivalent to (possibly trivial) wedges of m-spheres, and for $h(v)$ arbitrarily small are non-trivial wedges of m-spheres, then these hypotheses are satisfied.

3. (Pure) braid groups

Let B_n denote the n-strand braid group. The *standard presentation* for B_n is

$$B_n = \langle s_1, \ldots, s_{n-1} \mid s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1} \text{ for } 1 \leq i \leq n - 2$$
$$\text{and } s_i s_j = s_j s_i \text{ for } |i - j| > 1 \rangle.$$

Pictorially, s_i is the braid in which the ith strand crosses in front of the $(i+1)$st strand. This is called a *positive* crossing, and the reverse is a *negative* crossing. By imposing the additional relations $s_i^2 = 1$ for all i, one obtains the *standard presentation* for the symmetric group S_n. Hence there is an epimorphism $\pi \colon B_n \to S_n$ and its kernel is the n-strand *pure braid group* P_n. We take the action of S_n on $\{1, \ldots, n\}$ to be a right action, so the notation $(i)\sigma$ will be common. We will always label the strands of a braid at their tops, by the numbers 1 through n, from left to right. If we then labeled the strands at the bottom, 1 through n left to right, then each strand would be labeled i at the top and $(i)\pi(x)$ at the bottom for some i.

An element of B_n that will be important in all that follows is

$$\Delta := s_1 \cdots s_{n-1} s_1 \cdots s_{n-2} \cdots s_1 s_2 s_1.$$

This is the element in which each strand crosses in front of every strand to its right, exactly once. See Figure 1. Note that every crossing in Δ is positive.
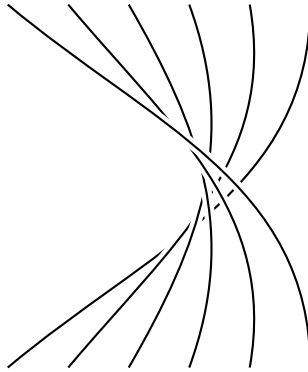

FIGURE 1. The element Δ in B_6.

We collect here some well known facts about Δ. These can be found for example in [**KT**, Section 1.3.3 and Chapter 6].

Citation 3.1 (Δ facts)**.** *The permutation $\pi(\Delta) \in S_n$ is the so called longest word w_0, which interchanges i and $n-i+1$, for all $1 \leq i \leq \lfloor n/2 \rfloor$. Using the defining relations, one can see that there exists, for any $1 \leq i \leq n-1$, a minimal word representing Δ and beginning with s_i. Also, it is clear that $s_i\Delta = \Delta s_{n-i}$ for all $1 \leq i \leq n-1$. For $n \geq 3$, the center of B_n is infinite cyclic, generated by Δ^2. Since $w_0^2 = \mathrm{id}$, the center lies in P_n, and in fact the center of P_n also equals $\langle \Delta^2 \rangle$ (for all n).*

Definition 3.2 (Automorphisms)**.** There are some useful automorphisms of P_n that we discuss now. First we have the conjugation action of B_n on P_n. Later we will consider characters $\omega_{i,j}$ of P_n, indexed by the labeling of the strands, and the conjugation action of B_n will permute the indices i, j. The other important automorphism is the inversion automorphism. This takes a braid diagram and changes all the right-over-left crossings to left-over-right, and vice versa. In particular if the automorphism is denoted μ, we have $\mu(s_i) = s_i^{-1}$ for all i, and $\mu(\Delta) = \Delta^{-1}$. Note that μ does not change the labeling of the strands.

Definition 3.3 (Natural projections)**.** Following [**KMM**], we define for any $A \subseteq \{1, \ldots, n\}$ an epimorphism $\phi_A \colon P_n \to P_{|A|}$ given by erasing all the strands labeled by numbers not in A. Here as before we always label strands 1 through n from left to right. We call the ϕ_A *natural projections*. They are in fact split epimorphisms, with the splitting given by embedding $P_{|A|}$ into P_n by adding the missing strands in the back. See also [**KT**, Section 1.3.2] for more details.

3.1. Characters. It is easy to abelianize B_n using the standard presentation, and find that $B_n^{ab} \cong \mathbb{Z}$. Hence we have $\mathrm{Hom}(B_n, \mathbb{R}) \cong \mathbb{R}$, and it is noteworthy that this is independent of n. In contrast, one can find presentations for P_n (see for example [**MM**]) that reveal that $P_n^{ab} \cong \mathbb{Z}^{\binom{n}{2}}$. Hence $\mathrm{Hom}(P_n, \mathbb{R}) \cong \mathbb{R}^{\binom{n}{2}}$, the dimension of which grows quickly with n.

There are bases of these vector spaces that correspond to natural measurements on braids. For B_n, a nice basis character, which we will call κ, is the one that reads off the total number of positive crossings of the strands, minus the total number of negative crossings. The character κ is induced by sending every standard generator of B_n to 1.

Corollary 3.4. $\Sigma^\infty(B_n) = S(B_n)$.

Proof: The points of $S(B_n)$ are $[\kappa]$ and $[-\kappa]$. Since $\kappa(\Delta^2) \neq 0$, the result follows from Corollary 1.6. $\square$

In particular Citation 1.2 implies that every coabelian subgroup of B_n (for example $[B_n, B_n]$) is of type F_∞.

What we actually care about is the pure braid groups; the character sphere $S(P_n)$ has dimension $\binom{n}{2} - 1$, so here the situation is much more complicated. A standard basis for the vector space $\mathrm{Hom}(P_n, \mathbb{R})$ is given by, for each $1 \leq i < j \leq n$, the *winding number characters* $\omega_{i,j}$. The character $\omega_{i,j} \colon P_n \to \mathbb{Z}$ is defined to be

$$\omega_{i,j} := \theta \circ \phi_{\{i,j\}},$$

where $\phi_{\{i,j\}}$ is the natural projection $P_n \to P_2$ and $\theta\colon P_2 \to \mathbb{Z}$ is the isomorphism sending Δ^2 to 1. Pictorially, $\omega_{i,j}\colon P_n \to \mathbb{Z}$ reads off the number of times the ith and jth strands wind completely around each other. The character $\omega_{i,j}$ can also be denoted $\omega_{j,i}$.

The function $\omega_{i,j}$ makes sense on B_n too, it just is not a homomorphism. Let $\zeta\colon B_2 \to \frac{1}{2}\mathbb{Z}$ be the isomorphism sending Δ to $\frac{1}{2}$, so ζ restricted to P_2 is θ. For $A \subseteq \{1,\ldots,n\}$ let $\psi_A\colon B_n \to B_{|A|}$ be the function (*not* homomorphism) that takes a braid in B_n, with strands labeled 1 through n at the top, and erases all the strands except those labeled by elements of A. Now define

$$\omega_{i,j} := \zeta \circ \psi_{\{i,j\}}.$$

We will use the notation $\omega_{i,j}$ for both the function $B_n \to \frac{1}{2}\mathbb{Z}$ and the homomorphism $P_n \to \mathbb{Z}$ since they coincide on P_n. We emphasize that when computing $\omega_{i,j}$ of a braid in B_n, we label the strands at the top, not the bottom. Note that for any $x \in B_n$ we have

$$\sum_{1 \le i < j \le n} 2\omega_{i,j}(x) = \kappa(x).$$

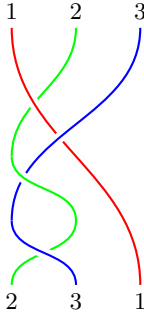See Figure 2 for some examples of $\omega_{i,j}$-values on a non-pure braid.

FIGURE 2. We label the strands at the top to compute the $\omega_{i,j}$. The different colors for the strands and the labeling at the bottom is just to make it easier to see how the strands interact. We compute that $\omega_{1,2} = \frac{1}{2}$, $\omega_{1,3} = -\frac{1}{2}$, and $\omega_{2,3} = 1$.

Despite $\omega_{i,j}\colon B_n \to \frac{1}{2}\mathbb{Z}$ not being a homomorphism, we do have the following:

Observation 3.5. *Let $x, y \in B_n$ and $1 \leq i < j \leq n$. Then*

$$\omega_{i,j}(xy) = \omega_{i,j}(x) + \omega_{(i)\pi(x),(j)\pi(x)}(y).$$

Proof: This is just a matter of keeping track of the labeling of the strands. The strand labeled i at the top of xy is labeled $(i)\pi(x)$ at the top of y, and similarly for j. Hence the total winding number of strands i and j is their winding number through x, plus their winding number through y, and we get the desired equation. See Figure 3 for an example. □
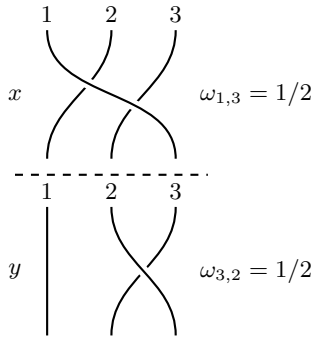

FIGURE 3. Here $\omega_{1,3}(xy) = 1$, $\omega_{1,3}(x) = 1/2$, and $\omega_{(1)\pi(x),(3)\pi(x)}(y) = \omega_{3,2}(y) = 1/2$, so the equation in Observation 3.5 is satisfied.

Taking linear combinations of the $\omega_{i,j}$, we can view any character χ of P_n as a function $B_n \to \mathbb{R}$. If $\chi = \sum a_{i,j}\omega_{i,j}$, denote by χ^x the character

$$\chi^x := \sum_{1 \leq i < j \leq n} a_{i,j}\omega_{(i)\pi(x),(j)\pi(x)},$$

so $\chi(xy) = \chi(x) + \chi^x(y)$ for all $x, y \in B_n$.

Before this subsection we discussed automorphisms of P_n. We now know how conjugation by elements of B_n affects characters, namely conjugation by x takes χ to χ^x (this is the reason for the notation). The other important automorphism was μ, which is induced by inverting each generator. This takes any character χ to $-\chi$. In particular since the BNSR-invariants are invariant under automorphisms, we now know that they are closed under taking antipodes:

Observation 3.6 (Antipodes). *If $[\chi] \in \Sigma^m(P_n)$ then $[-\chi] \in \Sigma^m(P_n)$.*

We now see more evidence that the element Δ is important.

Lemma 3.7. *Let $\chi\colon P_n \to \mathbb{R}$ be any character, and view χ as a function from B_n to $\mathbb{R}$ as above. Then for any $x \in B_n$, we have $\chi(x\Delta) = \chi(x) + \chi(\Delta)$.*

Proof: It suffices to check this for $\chi = \omega_{i,j}$. By Observation 3.5, we just need to prove that $\omega_{(i)\pi(x),(j)\pi(x)}(\Delta) = \omega_{i,j}(\Delta)$ for all $x \in B_n$ and all i, j. Indeed, $\omega_{i,j}(\Delta) = 1/2$ for all i and j, so we are done.   □

Corollary 3.8 (Δ survives). *Let $\chi\colon P_n \to \mathbb{R}$ be any character, and view χ as a function from B_n to $\mathbb{R}$ as above. If $\chi(\Delta) \neq 0$ then $[\chi] \in \Sigma^\infty(P_n)$.*

Proof: By Lemma 3.7, $\chi(\Delta^2) = 2\chi(\Delta)$, so if this is non-zero then $\chi(Z(P_n)) \neq 0$, and we are done by Corollary 1.6.   □

Since P_n has an $(n-1)$-dimensional compact classifying space, for example the Brady complex [**Bra**], Citation 1.5 tells us that $\Sigma^{n-1}(P_n) = \Sigma^\infty(P_n)$. We can do one better than this though:

Theorem 3.9. *For any $n \in \mathbb{N}$, $\Sigma^{n-2}(P_n) = \Sigma^\infty(P_n)$.*

Proof: First note that the character $\omega_{1,2}\colon P_n \to \mathbb{Z}$ splits, via the map $\mathbb{Z} \to P_n$ sending 1 to Δ^2, and Δ^2 is central, so $P_n \cong (P_n/Z(P_n)) \times \mathbb{Z}$. (Of course we could have used any $\omega_{i,j}$.) Now, $H := P_n/Z(P_n)$ is isomorphic to the mapping class group of the $(n+1)$-punctured sphere (see for example [**FM**, Chapter 9]), which has cohomological dimension $n-2$ [**Har**] and acts freely cocompactly on a contractible $(n-2)$-dimensional complex (see for example [**AM**, Corollary 1.3] for an even stronger result). By Citation 1.5 then, $\Sigma^{n-2}(H) = \Sigma^\infty(H)$.

We now return to $P_n \cong H \times \mathbb{Z}$ itself. Let $\chi\colon P_n \to \mathbb{R}$ be a non-trivial character, and suppose $[\chi] \in \Sigma^{n-2}(P_n)$. If $\chi(Z(P_n)) \neq 0$ then $[\chi] \in \Sigma^\infty(P_n)$ by Corollary 1.6, so assume $\chi(Z(P_n)) = 0$. Then χ is induced via $P_n \to H$ by a character χ' of H. By Citation 1.3, $[\chi'] \in \Sigma^{n-2}(H)$. As we have just seen, this means $[\chi'] \in \Sigma^\infty(H)$. Now since H is of type F_∞ and P_n is a trivial HNN-extension of H, we conclude, for example from [**BGK**, Theorem 2.3], that $[\chi] \in \Sigma^\infty(P_n)$.   □

4. The complex

In this section we define a locally compact simplicial complex X on which the braid group B_n acts freely and vertex transitively. Since P_n has finite index in B_n, its action on X is cocompact, and hence can be used to "reveal" the BNSR-invariants of P_n. This complex arises from a *Garside structure* on the symmetric group S_n, and has appeared

before, see for example [**CMW**]. For our purposes it will be useful to view X in a particular, slightly unconventional way, so we will start at the beginning and stay mostly self-contained.

Recall the standard group presentation for the braid group B_n:

$$\langle s_1, \ldots, s_{n-1} \mid s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1} \text{ for } 1 \leq i \leq n - 2$$
$$\text{and } s_i s_j = s_j s_i \text{ for } |i - j| > 1 \rangle.$$

Viewing this as a monoid presentation, we get the monoid B_n^+ of positive braids, that is, the subset of B_n consisting of elements representable by words in the s_i (and not the s_i^{-1}). Pictorially, a *positive braid* is one representable in such a way that every crossing is positive. Note that $\kappa \colon B_n^+ \to \mathbb{N}_0$ is a monoid homomorphism, and $\kappa(p) = 0$ for $p \in B_n^+$ if and only if $p = 1$, so B_n^+ has no non-trivial units.

Define a relation $\leq$ on B_n by saying $x \leq y$ whenever $xp = y$ for $p \in B_n^+$. Note that left multiplication by B_n preserves this relation. Since B_n^+ is a monoid with no non-trivial units, $\leq$ is a partial order. Since $s_i \leq \Delta$ for all i and Δ^2 is central (Citation 3.1), every $x \in B_n$ satisfies $x \leq \Delta^k$ for some $k > 0$. Thus, any two elements of B_n have a common upper bound, i.e. the poset is directed. In particular, the geometric realization $|B_n|$ is a contractible space on which B_n acts freely (the action is free on vertices by construction, and it is easy to see that the stabilizer of a vertex equals the intersection of its vertex stabilizers). It is not, however, cocompact; indeed it is not even finite dimensional. Our next goal is to retract $|B_n|$ down to a B_n-invariant subcomplex that is cocompact.

Definition 4.1. For $x \leq y$, write $x \preceq y$ if moreover $y \leq x\Delta$.

These next two results follow quickly from the definition.

Lemma 4.2. *If $x \preceq y$ then $y \preceq x\Delta$ and $y\Delta^{-1} \preceq x$.*

Proof: We are told that $x \leq y \leq x\Delta$. Say $y = xp$ for $p \in B_n^+$. Then since $s_i \Delta = \Delta s_{n-i}$ for all $1 \leq i \leq n-1$ (see Citation 3.1), we see that $\Delta^{-1} p \Delta \in B_n^+$. In particular $y\Delta = x\Delta(\Delta^{-1} p \Delta)$ tells us that $x\Delta \leq y\Delta$, and so indeed $y \preceq x\Delta$. Now say $x\Delta = yq$ for $q \in B_n^+$. Then since $\Delta q \Delta^{-1} \in B_n^+$, similar to the previous case, the equation $y\Delta^{-1}(\Delta q \Delta^{-1}) = x$ tells us that $y\Delta^{-1} \leq x$, and so $y\Delta^{-1} \preceq x$. $\square$

Observation 4.3. *If $x \leq y \leq z$ and $x \preceq z$ then $x \preceq y \preceq z$.*

Proof: We are told that $z \leq x\Delta$. Since $y \leq z$ we immediately see that $y \leq x\Delta$, so $x \preceq y$. Now we claim that $x\Delta \leq y\Delta$, after which we will get $z \leq y\Delta$ and hence $y \preceq z$. Choose $p \in B_n^+$ such that $xp = y$. As in the

previous proof, $\Delta^{-1}p\Delta \in B_n^+$. Hence $x\Delta(\Delta^{-1}p\Delta) = y\Delta$ tells us that indeed $x\Delta \leq y\Delta$. □

We now let X be the subcomplex of $|B_n|$ consisting of those simplices, i.e. chains $x_0 < \cdots < x_k$, such that $x_0 \preceq x_k$. By Observation 4.3, for such a simplex we have $x_i \preceq x_j$ for all i and j, so this property is closed under passing to faces and X is indeed a subcomplex. Using interval notation in the poset $(B_n, \leq)$, every simplex of X lies in the realization of an interval of the form $[x, x\Delta]$, and any such interval is finite, so X is locally compact. Also note that $X^{(0)} = |B_n|^{(0)}$, so B_n acts transitively on $X^{(0)}$, which means that $B_n \backslash X$ is compact.

Proposition 4.4. *X is contractible.*

Proof: Any simplex of $|B_n|$ lies in the realization of an interval of the form $|[x, xp]|$ for $x \in B_n$ and $p \in B_n^+$. If this realization is not contained in X, then Observation 4.3 says $x \npreceq xp$. Hence we can build up from X to $|B_n|$ by gluing in realizations of such intervals along their relative links, and we claim we can do this in such a way that said relative links are always contractible. This will then imply that the homotopy type never changes, and indeed $X \simeq |B_n|$ is contractible.

We glue in the intervals $|[x, xp]|$ for $x \npreceq xp$ in order of increasing $\kappa(p)$ value. Hence, when we attach $|[x, xp]|$, we do so along a relative link equal to $|[x, xp)| \cup |(x, xp]|$. This is just the suspension of $|(x, xp)|$, so it suffices to prove that for $x \npreceq xp$, $|(x, xp)|$ is contractible. To clean up notation, without loss of generality $x = 1$, so our assumption that $x \npreceq xp$ becomes $1 \npreceq p$, which just means that $p \nleq \Delta$.

Any $q \in B_n^+$ has a unique greatest lower bound with Δ, denoted $q \wedge \Delta$ (see for example [**KT**, Theorems 6.19 and 6.20]). Let $f \colon (1, p) \to (1, p)$ be the map $q \mapsto q \wedge \Delta$. It is clear that $f(q) \in [1, p)$ for $q \in (1, p)$, and it is less clear but still true that it is in $(1, p)$. This is because $q \in B_n^+ \backslash \{1\}$, and any such q satisfies $1 < q \wedge \Delta$ since Δ admits minimal representations beginning with any positive generator (see Citation 3.1). Since $f(q) \leq q$ and any element $f(q)$ in the image of f satisfies $f(q) \leq p \wedge \Delta$, if we can show that $p \wedge \Delta \in (1, p)$ then we will have a conical contraction of $(1, p)$ to a point, by [**Qui**, Section 1.5]. Here is where we use the fact that $p \nleq \Delta$; this implies that $p \wedge \Delta \neq p$, so indeed $p \wedge \Delta \in (1, p)$ and we are done. □

Remark 4.5. There are various other useful models for classifying spaces of braid groups, e.g., the Brady complex [**Bra**]. The Brady complex is similar to X, but with smaller lattices than our $[1, \Delta]$, namely *non-crossing partition lattices*. The Brady complex has many advantages,

for example it has the smallest dimension possible, but for our purposes X is the most germane model, since once characters enter the picture, Δ becomes very useful.

4.1. Character height functions. Recall from Subsection 3.1 that we can view characters χ of P_n as functions on B_n, which might just not be homomorphisms. Now, $B_n = X^{(0)}$, so extending χ affinely to the simplices of X, we get a function $h_\chi \colon X \to \mathbb{R}$. To make the notation cleaner, from now on we will just write χ for both the character on P_n and the function on X. Indeed, h_χ of a vertex was already also called χ, so calling the entire thing χ is not even really an abuse of notation.

Observation 4.6. *Let χ be a character of P_n. Viewing χ as a function on X, it is a character height function.*

Proof: For a vertex $x \in X$ and a group element $g \in P_n$, since $\pi(g) = \mathrm{id}$, Observation 3.5 says $\chi(gx) = \chi(g) + \chi(x)$. $\square$

More generally, Observation 3.5 says $\chi(xy) = \chi(x) + \chi^x(y)$ for $x, y \in B_n$ (with notation explained after the proof of Observation 3.5). In particular, if x is a vertex of X (so really an element of B_n) and $p \in [\Delta^{-1}, 1) \cup (1, \Delta]$, so xp is an adjacent vertex to x, then $\chi(xp) = \chi(x) + \chi^x(p)$. We see that $\chi(xp) > \chi(x)$ if and only if $\chi^x(p) > 0$. Since $\chi^x(1) = 0$, this can be phrased as saying that xp is χ-ascending relative x if and only if p is χ^x-ascending relative 1.

4.2. The weak Bruhat lattice. The vertex links in X are important to understand, and they are related to the *weak Bruhat lattice* W_n of the symmetric group S_n. This section is devoted to W_n.

We will write elements σ of S_n as brackets $\sigma = [k_1, \ldots, k_n]$, where this denotes that, for each $1 \le i \le n$, σ sends i to j such that $k_j = i$. In a braid x, where we label the strands from 1 to n at the top, if we propagate each label to the bottom of its strand, and call the new order $k_1, \ldots, k_n$ from left to right, then we get $\pi(x) = [k_1, \ldots, k_n]$. For example, going back to the braid x from Figure 2, we have $\pi(x) = [2, 3, 1]$, which is the permutation taking 2 to 1, 3 to 2, and 1 to 3.

Definition 4.7 (Weak Bruhat order). Let $\sigma, \tau \in S_n$. Define a relation $\le$ by saying that $\sigma \le \tau$ whenever, for all $1 \le i < j \le n$, if $(i)\sigma > (j)\sigma$ then $(i)\tau > (j)\tau$ (recall we use right actions). This is clearly reflexive, antisymmetric and transitive, so $(S_n, \le)$ is a poset. We call $\le$ the *weak Bruhat order*. The poset $(S_n, \le)$ is in fact a lattice [**BEZ**] with minimum $[1, \ldots, n]$ and maximum $[n, \ldots, 1]$.

We denote the geometric realization $|(S_n, \leq)|$ by W_n, and call it the *weak Bruhat lattice*. The *proper part* of W_n, denoted PW_n, is the subcomplex supported on those vertices other than $[1, \ldots, n]$ and $[n, \ldots, 1]$. The reason we are so interested in the weak Bruhat order is the following:

Lemma 4.8. *The interval $[1, \Delta]$ in the poset $(B_n, \leq)$ is isomorphic as a poset to $(S_n, \leq)$, the symmetric group with the weak Bruhat order.*

Proof: Let $\pi \colon [1, \Delta] \to S_n$ be the restriction of the usual projection $B_n \to S_n$. This restriction is bijective, by [**KT**, Lemma 6.24], so we just need to show it is a poset map. Suppose $x \leq y$ in B_n. We need to show that $\pi(x) \leq \pi(y)$. Say $xp = y$, so for any $1 \leq i < j \leq n$ we have $\omega_{i,j}(y) = \omega_{i,j}(x) + \omega_{(i)\pi(x),(j)\pi(y)}(p)$ by Observation 3.5. If $(i)\pi(x) > (j)\pi(x)$ then since $1 \leq x \leq \Delta$, we know $\omega_{i,j}(x) = 1/2$. Also, since $1 \leq y \leq \Delta$ we know $\omega_{i,j}(y)$ is either 0 or $1/2$, so since $\omega_{(i)\pi(x),(j)\pi(y)}(p) \geq 0$ this implies $\omega_{i,j}(y) = 1/2$. Hence $(i)\pi(y) > (j)\pi(y)$ and indeed $\pi(x) \leq \pi(y)$. $\square$

In the future we may suppress this isomorphism, and use the language of the weak Bruhat order when talking about the interval $[1, \Delta]$. Since we are using $\leq$ to denote the order in both places anyway, this may even pass unnoticed.

There is a family of subcomplexes of W_n that will be important later. For $1 \leq i < j \leq n$, the *proper (i,j)-reversing subcomplex* $\mathrm{Rev}_n(i,j)$ is the subcomplex of PW_n supported on those vertices σ with $(i)\sigma > (j)\sigma$. In bracket notation, if $\sigma = [k_1, \ldots, k_n]$ then $\sigma \in \mathrm{Rev}_n(i,j)$ if and only if j comes before i in the list $k_1, \ldots, k_n$. If we allowed $[n, \ldots, 1]$, this would be a cone point for $\mathrm{Rev}_n(i,j)$, but since $\mathrm{Rev}_n(i,j)$ lies in PW_n, computing its homotopy type requires some work.

Lemma 4.9. *For $1 \leq i < j \leq n$ with $1 < i$ or $j < n$, $\mathrm{Rev}_n(i,j)$ is contractible.*

Proof: We will cover $\mathrm{Rev}_n(i,j)$ by contractible subcomplexes, namely stars of certain vertices, and then use the Nerve Lemma [**BLVŽ**, Lemma 1.2]. Call a vertex of $\mathrm{Rev}_n(i,j)$ *minimal* if it is minimal under the ordering of PW_n restricted to $\mathrm{Rev}_n(i,j)$. Clearly any simplex of $\mathrm{Rev}_n(i,j)$ lies in the star of some minimal vertex. Note that if $x \in \mathrm{Rev}_n(i,j)$ and $x \leq y$ for $y \in PW_n$, then also $y \in \mathrm{Rev}_n(i,j)$. In particular, $\mathrm{Rev}_n(i,j)$ is closed under taking upper bounds not equal to $[n, \ldots, 1]$. Thanks to this, if a collection of stars of minimal vertices has a non-empty intersection, then their intersection is itself a star, namely the star of the join of these vertices. In particular, the Nerve Lemma says that $\mathrm{Rev}_n(i,j)$ is homotopy equivalent to the nerve of the covering by stars of minimal vertices.

Note that σ is minimal if and only if $\sigma = [k_1, \ldots, k_r, j, i, k_{r+3}, \ldots, k_n]$ for some $k_1 < \cdots < k_r < j$ and $i < k_{r+3} < \cdots < k_n$. Hence, the join in W_n of all the minimal vertices of $\mathrm{Rev}_n(i, j)$ is $[1, \ldots, i-1, j, j-1, \ldots, i+1, i, j+1, \ldots, n]$. Since $1 < i$ or $j < n$, this lies in PW_n and hence in $\mathrm{Rev}_n(i, j)$. In particular, the nerve of the covering is a simplex, implying $\mathrm{Rev}_n(i, j)$ is contractible. $\qquad\square$

Lemma 4.10. *For $n \geq 3$, $\mathrm{Rev}_n(1, n)$ is homotopy equivalent to an $(n-3)$-sphere.*

Proof: As in the previous proof, we will cover $\mathrm{Rev}_n(1, n)$ by stars of certain vertices and use the Nerve Lemma. This time we want to use *maximal* vertices. The maximal vertices of PW_n are those of the form $[n, \ldots, i+2, i, i+1, i-1, \ldots, 1]$. That is, for each $1 \leq i < n$ we have a maximal vertex v_i in which every pair of entries is reversed except for $(i, i+1)$. Since $n \geq 3$, the pair $(1, n)$ is not of this form, so all these maximal vertices of PW_n lie in $\mathrm{Rev}_n(1, n)$. Clearly any simplex of $\mathrm{Rev}_n(1, n)$ lies in the star of some maximal vertex, so $\mathrm{Rev}_n(1, n)$ is covered by the stars of its maximal vertices, of which there are $n-1$. If a collection of such stars has non-empty intersection in $\mathrm{Rev}_n(1, n)$, then it equals the star of the meet of the relevant maximal vertices. Hence by the Nerve Lemma, $\mathrm{Rev}_n(1, n)$ is homotopy equivalent to the nerve of this covering.

Now, given any $n-2$ maximal vertices, say all of them except v_i, the meet of these is $[i+1, \ldots, n, 1, \ldots, i]$. Since this lies in $\mathrm{Rev}_n(1, n)$, every proper subset of vertices of the nerve spans a simplex. However, the collection of all the vertices does not span a simplex, since the meet in W_n of all the maximal vertices is $[1, \ldots, n]$, and this does not lie in $\mathrm{Rev}_n(1, n)$, in fact it does not even lie in PW_n. Hence the nerve is the boundary of an $(n-2)$-simplex, which is homotopy equivalent to an $(n-3)$-sphere. $\qquad\square$

5. Main results

Since the action of P_n on X is free and cocompact, and X is contractible, for any $[\chi] \in S(P_n)$ we know that $[\chi] \in \Sigma^m(P_n)$ if and only if the filtration $(X_{\chi \geq t})_{t \in \mathbb{R}}$ is essentially $(m-1)$-connected. We now encode χ into a Morse function as follows. Consider the following total ordering $\lessdot$ on the integers:

$$\cdots \lessdot 3 \lessdot 2 \lessdot 1 \lessdot \cdots \lessdot -3 \lessdot -2 \lessdot -1 \lessdot 0.$$

In words, the non-positive numbers are ordered like normal, but the positive numbers are in the reverse order, and are all less than every

non-positive number. Recall the "number of crossings" homomorphism $\kappa\colon B_n \to \mathbb{Z}$. We will write $\dot{\kappa}$ for the function κ with its outputs ordered according to $\lessdot$. As functions of sets, κ and $\dot{\kappa}$ are identical, the dot is just to emphasize the unusual ordering of the outputs. We will also write $\dot{\mathbb{Z}}$ for the set $\mathbb{Z}$ with the ordering $\lessdot$.

Extend $\dot{\kappa}$ to a map $X \to \mathbb{R}$ by affinely extending it to the simplices. Here we have fixed some order-preserving embedding of $\dot{\mathbb{Z}}$ into $\mathbb{R}$, but it will not matter what it is precisely. Consider the lexicographically ordered function

$$(\chi, \dot{\kappa})\colon X \to \mathbb{R} \times \mathbb{R}$$

on X.

Lemma 5.1. $(\chi, \dot{\kappa})$ *is an ascending-type Morse function.*

Proof: By construction, χ and $\dot{\kappa}$ are affine on cells. Since P_n acts cocompactly, $\chi(v) - \chi(w)$ takes only finitely many values on pairs of adjacent vertices (v, w), so we can choose ε to be less than all of the non-zero such values. Then since $\dot{\kappa}$ always takes different values on adjacent vertices, and since the outputs of $\dot{\kappa}$ are inversely well ordered by construction, indeed $(\chi, \dot{\kappa})$ is an ascending-type Morse function.   □

For a vertex x and a neighboring vertex y, y is in the ascending link of x if and only if either $\chi(y) > \chi(x)$, or else $\chi(y) = \chi(x)$ and $\dot{\kappa}(y) \gtrdot \dot{\kappa}(x)$. If $\kappa(x)$ is positive (that is, $\kappa(x) > 0$ in the usual ordering), then $\dot{\kappa}(y) \gtrdot \dot{\kappa}(x)$ is equivalent to $\kappa(y) < \kappa(x)$. If $\kappa(x)$ is non-positive, then $\dot{\kappa}(y) \gtrdot \dot{\kappa}(x)$ is equivalent to $0 \geq \kappa(y) > \kappa(x)$.

Recall from Corollary 3.8 that if $\chi(\Delta) \neq 0$ then $[\chi] \in \Sigma^\infty(P_n)$, so this case is finished. For the rest of the section, we will assume that $\chi(\Delta) = 0$. In particular, in this case Lemma 3.7 tells us that

(5.1) $\chi(x\Delta) = \chi(x)$

for all $x \in B_n$.

Lemma 5.2. *If $\kappa(x)$ is positive then* lk$\uparrow x$ *is homotopy equivalent to* lk$\uparrow x \cap [x\Delta^{-1}, x)$, *which is contractible.*

Proof: Since $\kappa(x)$ is positive and $\kappa(x\Delta^{-1}) < \kappa(x)$, we get that $\dot{\kappa}(x\Delta^{-1}) \gtrdot \dot{\kappa}(x)$. Also, $\chi(x\Delta^{-1}) = \chi(x)$ by Equation (5.1), so $x\Delta^{-1} \in$ lk$\uparrow x$. Hence the intersection lk$\uparrow x \cap [x\Delta^{-1}, x)$ is a contractible cone on $x\Delta^{-1}$.

Now we need to show the homotopy equivalence. Let $L :=$ lk$\uparrow x$ and $N :=$ lk$\uparrow x \cap [x\Delta^{-1}, x)$. Since L is finite and κ takes distinct values on adjacent vertices, κ is a (classical) Morse function on L (here we use the usual ordering, i.e. we are using κ and not $\dot{\kappa}$). By definition,

N is the sublevel set $L_{\kappa<\kappa(x)}$. Hence by Corollary 2.5, it suffices to show that for any vertex y of L with $\kappa(y) > \kappa(x)$, the κ-descending link $\operatorname{lk}_L^{\kappa\downarrow} y$ of y in L is contractible. For such a y, we claim that $\operatorname{lk}_L^{\kappa\downarrow} y$ is a contractible cone on $y\Delta^{-1}$. First we need to check that $y\Delta^{-1}$ actually lies in $\operatorname{lk}_L^{\kappa\downarrow} y$. Since $y \in L$ and $\kappa(y) > \kappa(x)$, we have $x \prec y$, so $y\Delta^{-1} \prec x$ by Lemma 4.2. This shows $y\Delta^{-1} \in \operatorname{lk}_X x$, but we also need it to be $(\chi, \dot\kappa)$-ascending. For this note that, since $\kappa(y) > \kappa(x) > 0$, we know $\dot\kappa(x) \gtrdot \dot\kappa(y)$, and so for y to be in L we must have $\chi(y) > \chi(x)$, and hence $\chi(y\Delta^{-1}) > \chi(x)$. We conclude that $y\Delta^{-1} \in L$. Since $\kappa(y\Delta^{-1}) < \kappa(y)$, we have $y\Delta^{-1} \in \operatorname{lk}_L^{\kappa\downarrow} y$. Now we claim any $z \in \operatorname{lk}_L^{\kappa\downarrow} y$ satisfies $y\Delta^{-1} \preceq z$. Indeed, such a z satisfies $z \prec y$, so this follows from Lemma 4.2. We conclude that $\operatorname{lk}_L^{\kappa\downarrow} y$ is a contractible cone on $y\Delta^{-1}$. $\square$

Lemma 5.3. *If $\kappa(x)$ is non-positive then $\operatorname{lk}{\uparrow}x$ is homotopy equivalent to $\operatorname{lk}{\uparrow}x \cap (x, x\Delta]$.*

Proof: This follows by a somewhat parallel argument to the previous proof. Let $L := \operatorname{lk}{\uparrow}x$ and let $P := \operatorname{lk}{\uparrow}x \cap (x, x\Delta]$. Consider κ as a classical Morse function on L, so P is the superlevel set $L_{\kappa>\kappa(x)}$. We need to show that for any vertex y of L with $\kappa(y) < \kappa(x)$, the κ-ascending link $\operatorname{lk}_L^{\kappa\uparrow} y$ of y in L is contractible. We claim it is a cone on $y\Delta$. First, $y \prec x$ so $x \preceq y\Delta$ and $y\Delta \in \operatorname{lk}_X x$. Moreover, since $\kappa(y) < \kappa(x) \le 0$, we have $\dot\kappa(y) \lessdot \dot\kappa(x)$, so for y to lie in L we must have $\chi(y) > \chi(x)$. Hence also $\chi(y\Delta) > \chi(x)$ and so $y\Delta \in \operatorname{lk}_L^{\kappa\uparrow} y$. Lastly, for any $z \in \operatorname{lk}_L^{\kappa\uparrow} y$ we have $y \prec z$, so $z \preceq y\Delta$, and indeed $\operatorname{lk}_L^{\kappa\uparrow} y$ is a cone on $y\Delta$, hence contractible. $\square$

There is a conspicuous difference between these two lemmas, which is worth pointing out explicitly, namely, when $\kappa(x)$ is non-positive, $\operatorname{lk}{\uparrow}x \cap (x, x\Delta]$ could fail to be contractible. The asymmetry comes from the conditions to be ascending: when $\kappa(x)$ is positive, a vertex $y \in \operatorname{lk}_X x$ is ascending if either $\chi(y) > \chi(x)$, or $\kappa(y) < \kappa(x)$, whereas when $\kappa(x)$ is non-positive, y is ascending if either $\chi(y) > \chi(x)$, or $0 \ge \kappa(y) > \kappa(x)$. Hence, where $x\Delta^{-1}$ served as a cone point for $\operatorname{lk}{\uparrow}x \cap [x\Delta^{-1}, x)$, now in $\operatorname{lk}{\uparrow}x \cap (x, x\Delta]$ we might not even have $x\Delta$ to use as a cone point, e.g., if $\kappa(x\Delta) > 0$.

Let us give the complex P from the last proof a more official name.

Definition 5.4 (Positive ascending link). Define the *positive ascending link* $\operatorname{plk}{\uparrow}x$ of x to be $\operatorname{plk}{\uparrow}x := \operatorname{lk}{\uparrow}x \cap (x, x\Delta]$.

Corollary 5.5. *If $\kappa(x\Delta) \le 0$ then $\operatorname{lk}{\uparrow}x$ is contractible.*

Proof: Since $\kappa(x) \leq \kappa(x\Delta)$, we know $\kappa(x) \leq 0$. In this case, Lemma 5.3 says lk↑$x \simeq$ plk↑x, so our goal is to prove that plk↑x is contractible. Since $\kappa(x\Delta) \leq 0$, and since $\chi(x\Delta) = \chi(x)$ by Equation (5.1), $\kappa(x\Delta) \in$ plk↑x. But this means plk↑x is a cone on $x\Delta$, and so is contractible.   □

Thanks to Lemma 5.2 and Corollary 5.5, the only time lk↑x might fail to be contractible is when $\kappa(x) \leq 0 < \kappa(x\Delta)$. In this case Lemma 5.3 says that lk↑$x \simeq$ plk↑x, which equals the "χ-ascending part" of $(x, x\Delta)$ together with some amount of its "χ-flat part" added.

Here is our main application of all this setup:

Proposition 5.6. *Let* $\chi = \sum a_{i,j}\omega_{i,j}$ *be a character of* P_n. *Suppose that* $\sum a_{i,j} = 0$, *so* $\chi(\Delta) = 0$, *and that either exactly one* $a_{i,j}$ *is positive, or exactly one* $a_{i,j}$ *is negative. Then* $[\chi] \notin \Sigma^{n-2}(P_n)$. *If moreover none of the* $a_{i,j}$ *are zero, then* $[\chi] \in \Sigma^{n-3}(P_n)$.

Proof: Any such χ is the limit of a sequence of characters of the same form, and such that no $a_{i,j}$ is zero. Since $S(P_n) \setminus \Sigma^{n-2}(P_n)$ is closed, we may assume without loss of generality that no $a_{i,j}$ is zero, and then both statements in the theorem will follow if we show that $[\chi] \in \Sigma^{n-3}(P_n) \setminus \Sigma^{n-2}(P_n)$. By Observation 3.6, we may assume without loss of generality that exactly one $a_{i,j}$ is positive (and the others are negative). Say the positive one is $a_{k,\ell}$.

We inspect ascending links lk↑x with respect to $(\chi, \dot\kappa)$. Thanks to our assumptions, no proper non-empty subset of $\{a_{i,j} \mid 1 \leq i < j \leq n\}$ sums to zero, so the only vertices of lk$_X x$ with the same χ-value as x are $x\Delta$ and $x\Delta^{-1}$. By Lemma 5.2 and Corollary 5.5, lk↑x will be contractible unless $\kappa(x) \leq 0 < \kappa(x\Delta)$, so we can assume these bounds hold. In this case, Lemma 5.3 says lk↑$x \simeq$ plk↑x, but we also have $\dot\kappa(x\Delta) \prec \dot\kappa(x)$, so plk↑$x$ is just the subcomplex of $(x, x\Delta)$ supported on those vertices y with $\chi(y) > \chi(x)$. Call this the χ-*ascending part* of $(x, x\Delta)$. (There is no "χ-flat part" thanks to our assumptions.)

For $y \in (x, x\Delta)$, say with $y = xp$ for $p \in B_n^+$, we have that $\chi(y) > \chi(x)$ if and only if $\chi^x(p) > 0$, as discussed in Subsection 4.1. Hence the χ-ascending part of $(x, x\Delta)$ is isomorphic to the χ^x-ascending part of $(1, \Delta)$. Since $\chi^x = \sum a_{i,j}\omega_{(i)\pi(x),(j)\pi(x)}$, the lone positive coefficient $a_{k,\ell}$ in χ^x is on $\omega_{(k)\pi(x),(\ell)\pi(x)}$. Let p and q denote $(k)\pi(x)$ and $(\ell)\pi(x)$, in whichever order gives us $p < q$. Thus a vertex of $(1, \Delta)$ is χ^x-ascending if and only if it is $\omega_{p,q}$-ascending. If we view $(1, \Delta)$ as the proper part PW$_n$ of the weak Bruhat lattice W$_n$, then this is precisely the subcomplex Rev$_n(p, q)$ of PW$_n$. By Lemma 4.9, this is contractible unless $p = 1$ and $q = n$, and then it is homotopy equivalent to S^{n-3}.

In particular, $\mathrm{lk}{\uparrow}x$ is always $(n-4)$-connected, so Corollary 2.5 says $[\chi] \in \Sigma^{n-3}(P_n)$.

Now, the case when $\mathrm{lk}{\uparrow}x \simeq S^{n-3}$ happens precisely for those x for which $\{(k)\pi(x), (\ell)\pi(x)\} = \{1, n\}$, and it is clear that for any $M \in \mathbb{R}$ there exists x with $\chi(x) < M$ such that $\{(k)\pi(x), (\ell)\pi(x)\} = \{1, n\}$. Also, $\widetilde{H}_{n-2}(\mathrm{lk}{\uparrow}x) = 0$ for all x, so Proposition 2.6 tells us that $[\chi] \notin \Sigma^{n-2}(P_n)$. $\qquad \square$

For $3 \leq m \leq n$, consider the character

$$\chi_n^m := \left(\sum_{\substack{1 \leq i < j \leq m \\ (i,j) \neq (1,2)}} \omega_{i,j} \right) - \left(\binom{m}{2} - 1 \right) \omega_{1,2}.$$

Note that sum of the coefficients in χ_n^m is zero, exactly one coefficient is negative, and if $m = n$ then none of the coefficients is zero. In particular Proposition 5.6 tells us that $[\chi_m^m] \in \Sigma^{m-3}(P_m) \setminus \Sigma^{m-2}(P_m)$ and $[\chi_n^m] \notin \Sigma^{n-2}(P_n)$. We can do better than this though. To prove our main theorem, that $\Sigma^{m-2}(P_n) \subseteq \Sigma^{m-3}(P_n)$ is a proper inclusion for all $3 \leq m \leq n$, we will prove that $[\chi_n^m] \in \Sigma^{m-3}(P_n) \setminus \Sigma^{m-2}(P_n)$.

Theorem 5.7 (Separation). *For any $3 \leq m \leq n$, the inclusion $\Sigma^{m-2}(P_n) \subseteq \Sigma^{m-3}(P_n)$ is proper. Explicitly, $[\chi_n^m] \in \Sigma^{m-3}(P_n) \setminus \Sigma^{m-2}(P_n)$.*

Proof: For the negative statement, note that $\chi_n^m = \chi_m^m \circ \phi_{\{1,\dots,m\}}$, so Citation 1.3 and Proposition 5.6 tell us that $[\chi_n^m] \notin \Sigma^{m-2}(P_n)$. For the positive statement, it is known that the kernel of $\phi_{\{1,\dots,n-1\}} \colon P_n \to P_{n-1}$ is isomorphic to the free group F_{n-1} [**KT**, Theorem 1.16], which is of type F_∞, so Lemma 1.4 and Observation 3.6 say that when a discrete character on P_n is induced from one on P_{n-1}, it inherits the latter's positive BNSR-invariant properties. By induction, when a discrete character on P_n is induced from one on P_m for any $m < n$, it inherits the latter's positive BNSR-invariant properties. Since $[\chi_m^m] \in \Sigma^{m-3}(P_m)$ by Proposition 5.6, we conclude that $[\chi_n^m] \in \Sigma^{m-3}(P_n) \setminus \Sigma^{m-2}(P_n)$. $\qquad \square$

Corollary 5.8. *For any $3 \leq m \leq n$, $\ker(\chi_n^m)$ is of type F_{m-3} but not F_{m-2}.*

Proof: This is immediate from Citation 1.2, Observation 3.6, and Theorem 5.7. $\qquad \square$

We have thus found examples of coabelian (even "cocyclic") subgroups of P_n with every possible finiteness length, namely 0 through $n - 2$.

Here is a nice, easy-to-state result in a related vein, which follows by combining our results with the full computation of $\Sigma^1(P_n)$ in [**KMM**].

Corollary 5.9. *For any $n \geq 4$, $\ker(\omega_{1,2} - \omega_{3,4}) \leq P_n$ is finitely generated but not finitely presentable.*

Proof: It suffices by Citation 1.2 and Observation 3.6 to prove that $[\omega_{1,2} - \omega_{3,4}] \in \Sigma^1(P_n) \setminus \Sigma^2(P_n)$. That it lies in Σ^1 follows from the complete computation of $\Sigma^1(P_n)$ done in [**KMM**]. Now we claim that $[\omega_{1,2} - \omega_{3,4}] \notin \Sigma^2(P_n)$. Using the natural projection $\phi_{\{1,2,3,4\}} \colon P_n \to P_4$, it suffices by Citation 1.3 to prove this in the $n = 4$ case, but this is immediate from Proposition 5.6. $\square$

Remark 5.10. Ideally one would like a complete computation of $\Sigma^m(P_n)$ for all m and n. Using our setup, if one can show that, for any $x \in B_n$ and any $0 \leq k < \kappa(\Delta)$, the subcomplex of PW_n supported on those vertices p with either $\chi^x(p) > 0$, or $\chi^x(p) = 0$ and $\kappa(p) \leq k$, is homotopy equivalent to a wedge of m-spheres, then we could conclude that $[\chi] \in \Sigma^{m-1}(P_n) \setminus \Sigma^m(P_n)$. However, for characters χ other than those of the type considered here, it is unclear at present how to analyze the homotopy types of these subcomplexes. Also we have found that this method will not always work, for example the Morse function $(\chi, \dot{\kappa})$ cannot fully recover the results of [**KMM**] on $\Sigma^1(P_n)$. Namely, for χ in a "P_4-circle", so $[\chi] \notin \Sigma^1(P_n)$, we have computed that there exist ascending links homotopy equivalent to S^0 and others homotopy equivalent to S^1, so Proposition 2.6 does not apply.

References

[AM] J. Aramayona and C. Martínez-Pérez, The proper geometric dimension of the mapping class group, *Algebr. Geom. Topol.* **14(1)** (2014), 217–227. DOI: 10.2140/agt.2014.14.217.

[BB] M. Bestvina and N. Brady, Morse theory and finiteness properties of groups, *Invent. Math.* **129(3)** (1997), 445–470. DOI: 10.1007/s002220050168.

[BGK] R. Bieri, R. Geoghegan, and D. H. Kochloukova, The sigma invariants of Thompson's group F, *Groups Geom. Dyn.* **4(2)** (2010), 263–273. DOI: 10.4171/GGD/83.

[BNS] R. Bieri, W. D. Neumann, and R. Strebel, A geometric invariant of discrete groups, *Invent. Math.* **90(3)** (1987), 451–477. DOI: 10.1007/BF01389175.

[BR]   R. BIERI AND B. RENZ, Valuations on free resolutions and higher geometric invariants of groups, *Comment. Math. Helv.* **63(1)** (1988), 464–497. DOI: 10.1007/BF02566775.

[BEZ]   A. BJÖRNER, P. H. EDELMAN, AND G. M. ZIEGLER, Hyperplane arrangements with a lattice of regions, *Discrete Comput. Geom.* **5(3)** (1990), 263–288. DOI: 10.1007/BF02187790.

[BLVŽ]   A. BJÖRNER, L. LOVÁSZ, S. T. VREĆICA, AND R. T. ŽIVALJEVIĆ, Chessboard complexes and matching complexes, *J. London Math. Soc. (2)* **49(1)** (1994), 25–39. DOI: 10.1112/jlms/49.1.25.

[Bra]   T. BRADY, A partial order on the symmetric group and new $K(\pi, 1)$'s for the braid groups, *Adv. Math.* **161(1)** (2001), 20–40. DOI: 10.1006/aima.2001.1986.

[BH]   M. R. BRIDSON AND A. HAEFLIGER, *"Metric Spaces of Non-Positive Curvature"*, Grundlehren der Mathematischen Wissenschaften **319**, Springer-Verlag, Berlin, 1999. DOI: 10.1007/978-3-662-12494-9.

[Bux]   K.-U. BUX, Finiteness properties of soluble arithmetic groups over global function fields, *Geom. Topol.* **8** (2004), 611–644. DOI: 10.2140/gt.2004.8.611.

[BG]   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(3)** (1999), 793–801. DOI: 10.1112/S0024610799007991.

[CMW]   R. CHARNEY, J. MEIER, AND K. WHITTLESEY, Bestvina's normal form complex and the homology of Garside groups, *Geom. Dedicata* **105(1)** (2004), 171–188. DOI: 10.1023/B:GEOM.0000024696.69357.73.

[FM]   B. FARB AND D. MARGALIT, *"A Primer on Mapping Class Groups"*, Princeton Mathematical Series **49**, Princeton University Press, Princeton, NJ, 2012.

[Har]   J. L. HARER, The virtual cohomological dimension of the mapping class group of an orientable surface, *Invent. Math.* **84(1)** (1986), 157–176. DOI: 10.1007/BF01388737.

[KT]   C. KASSEL AND V. TURAEV, *"Braid Groups"*, With the graphical assistance of Olivier Dodane, Graduate Texts in Mathematics **247**, Springer, New York, 2008. DOI: 10.1007/978-0-387-68548-9.

[KMM]   N. KOBAN, J. MCCAMMOND, AND J. MEIER, The BNS-invariant for the pure braid groups, *Groups Geom. Dyn.* **9(3)** (2015), 665–682. DOI: 10.4171/GGD/323.

[MM]   D. Margalit and J. McCammond, Geometric presentations for the pure braid group, *J. Knot Theory Ramifications* **18(1)** (2009), 1–20. DOI: 10.1142/S0218216509006859.

[MMV1] J. Meier, H. Meinert, and L. VanWyk, Higher generation subgroup sets and the Σ-invariants of graph groups, *Comment. Math. Helv.* **73(1)** (1998), 22–44. DOI: 10.1007/s000140050044.

[MMV2] J. Meier, H. Meinert, and L. VanWyk, On the Σ-invariants of Artin groups. Geometric topology and geometric group theory (Milwaukee, WI, 1997), *Topology Appl.* **110(1)** (2001), 71–81. DOI: 10.1016/S0166-8641(99)00170-4.

[Mei]   H. Meinert, Actions on 2-complexes and the homotopical invariant Σ^2 of a group, *J. Pure Appl. Algebra* **119(3)** (1997), 297–317. DOI: 10.1016/S0022-4049(96)00030-8.

[Qui]   D. Quillen, Homotopy properties of the poset of nontrivial p-subgroups of a group, *Adv. in Math.* **28(2)** (1978), 101–128. DOI: 10.1016/0001-8708(78)90058-0.

[WZ]   S. Witzel and M. C. B. Zaremsky, *"The Σ-Invariants of Thompson's Group F via Morse Theory"*, Topological Methods in Group Theory, London Math. Soc. Lecture Note Ser., Cambridge Univ. Press, Cambridge, 2015 (to appear).

[Zar]   M. C. B. Zaremsky, On the Σ-invariants of generalized Thompson groups and Houghton groups, *Int. Math. Res. Not. IMRN* (to appear). arXiv:1502.02620.

Department of Mathematical Sciences
Binghamton University
Binghamton, NY 13902
USA
E-mail address: zaremsky@math.binghamton.edu