AGT

Algebraic & Geometric Topology

Volume 25 (2025)

BNSR-invariants of surface Houghton groups

NOAH TORGERSON
JEREMY WEST

msp

msp

Algebraic & Geometric Topology **25**:8 (2025) 4897–4919
DOI: 10.2140/agt.2025.25.4897
Published: 20 November 2025

BNSR-invariants of surface Houghton groups

NOAH TORGERSON
JEREMY WEST

The surface Houghton groups $\mathcal{H}_n$ are a family of groups generalizing Houghton groups H_n, which are constructed as asymptotically rigid mapping class groups. We give a complete computation of the BNSR-invariants $\Sigma^m(P\mathcal{H}_n)$ of their intersection with the pure mapping class group. To do so, we prove that the associated Stein–Farley cube complex is CAT(0), and we adapt Zaremsky's method for computing the BNSR-invariants of the Houghton groups. As a consequence, we give a criterion for when subgroups of H_n and $P\mathcal{H}_n$ having the same finiteness length as their parent group are finite index. We also discuss the failure of some of these groups to be co-Hopfian.

20F65, 57K20, 57M07

1 Introduction

To any group G of type F_k, one can assign a sequence of invariants $\Sigma^m(G)$, for $m \leq k$. These invariants determine which subgroups of G (containing the commutator subgroup) share which finiteness properties of G. Historically, they are difficult to compute. They were defined across several papers (see [6; 7; 18]), primarily by Bieri, Neumann, Strebel, and Renz, hence the names "BNS-invariants" and "BNSR-invariants". For the remainder of this paper, these will be referred to as Σ-invariants for the sake of brevity. One recent collection of groups for which these have been computed is the family of Houghton groups [19; 20].

Houghton defined his groups as permutation groups of infinite sets in [15]. More specifically, H_n is the group of "eventual translations" of the set $\{1,\dots,n\}\times\mathbb{N}$, ie permutations which in each ray are translations outside some finite set. Brown proved, via a suitable simplicial complex, that Houghton's group H_n is of type F_{n-1}, but not of type F_n (in the language of "finiteness length" this is $\mathrm{fl}(H_n)=n-1$). (Brown proved type FP_{n-1} but not type FP_n, and finitely presented, which together imply F_n but not F_{n-1}.) In [2], the authors define a variant of Houghton groups, called surface Houghton groups, as asymptotically rigid mapping class groups of surfaces with infinite genus. We denote the surface Houghton groups by $\mathcal{H}_n$, and the pure (ie end-fixing) subgroups as $P\mathcal{H}_n$. In the same paper, they also show that $\mathrm{fl}(\mathcal{H}_n)=n-1$, via a cube complex analogous to the Stein–Farley complexes for Thompson groups.

This is not the only asymptotically rigid mapping class group variant of Houghton groups. While the surface Houghton groups replace the $\mathbb{N}$-rays with ends accumulated by genus (without boundary), the

© 2025 The Authors, under license to MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via Subscribe to Open.

braided Houghton groups (defined by Degenhardt in his PhD thesis [11]) can be realized as asymptotically rigid mapping class groups as well (see [12] for details). This construction replaces the $\mathbb{N}$-rays with planar ends, accumulated by punctures; this surface has noncompact boundary. The braided Houghton groups $\mathrm{Br}\, H_n$ also have $\mathrm{fl}(\mathrm{Br}\, H_n) = n-1$ (as proven in [14]). The braided Houghton groups shall play a small role in the final section of this paper.

In [19; 20], Zaremsky computed the Σ-invariants of the Houghton groups, using the cube complex defined implicitly in [8], and explicitly in [16]. We carry out a parallel of Zaremsky's arguments, showing that the equivalent statement holds for pure surface Houghton groups. Namely, we prove in Section 4 the following, where $m(\chi)$ is the number of nonzero coefficients of χ in ascending standard form (details in Section 4):

Theorem 4.1 *Let χ be a nonzero character of $P\mathcal{H}_n$. Then $[\chi] \in \Sigma^{m(\chi)-1}(P\mathcal{H}_n) \setminus \Sigma^{m(\chi)}(P\mathcal{H}_n)$.*

As an application of this, we provide a partial converse (for Houghton groups and surface Houghton groups) of the well-known fact that if $H \leq G$ is a finite-index subgroup, then $\mathrm{fl}(H) = \mathrm{fl}(G)$. Specifically, we show that whenever a subgroup of H_n or $P\mathcal{H}_n$ has finiteness length $n-1$ and intersects the corresponding commutator with finite index, it is finite index in the full group. Further, in order to be finite index, the intersection condition must hold, for general group-theoretic reasons. (These are also true of $\mathcal{H}_n$, but for trivial reasons: the commutator subgroup of $\mathcal{H}_n$ is all of $\mathcal{H}_n$.) This is carried out in Section 5, along with some discussion of co-Hopfianness.

In order to accomplish this, we demonstrate in Section 3 that the cube complex of [2] is CAT(0).

Theorem 3.9 *Let X_n denote the Stein–Farley complex for $\mathcal{H}_n$. Then X_n is* CAT(0).

To do this, we use a refinement of a proposition from [14]; see Proposition 3.2. In Section 2, we lay out the definitions of the surfaces and groups we are concerned with, of the corresponding cube complex, and of the version of discrete Morse theory we shall employ. Alongside these definitions are various lemmas which we shall need. We also obtain various nice representatives for the vertices (Section 2) and edges (Section 3) of the complex.

In Section 5, we finish with some applications of the Σ-invariants to when subgroups having maximal finiteness length have finite index. This discussion applies to H_n, $\mathcal{H}_n$, and $P\mathcal{H}_n$. To handle the infinite-index case, there is some discussion of the failure of these groups to be co-Hopfian. In particular, we have the following theorem.

Theorem 5.6 *Let H denote either the Houghton group, or the pure surface Houghton group, and suppose $G < H$ has $\mathrm{fl}(G) = \mathrm{fl}(H)$. Then G is finite index in H if and only if $G \cap H'$ is finite index in H', where H' denotes the commutator subgroup of H. Furthermore, there exist subgroups G with $\mathrm{fl}(G) = \mathrm{fl}(H)$ of both finite and infinite index.*

Recently, Marie Abadie [1] analyzed the CAT(0) property for the Stein–Farley cube complexes of a different family of asymptotically rigid mapping class groups, including the braided Houghton groups. Similar methods are used, in particular a version of the same proposition as we use. Thus, one could attempt to apply Zaremsky's methods to compute the Σ-invariants of Br H_n as well. It seems reasonable to guess that they should work out the same.

The recent paper [3] concerns a generalization of these surface Houghton groups, obtained by varying the rigid structure on the same surface. As they show that the groups they consider are finite-index subgroups of our $\mathcal{H}_n$, the Σ-invariants are effectively the same. Additionally, our results in Section 5 apply to these more general surface Houghton groups.

Acknowledgements

We would like to thank Jing Tao for advice and guidance on this project. We would also like to thank Noel Brady and Justin Malestein for valuable conversations, and Peter Patzt and Matt Zaremsky for helpful comments. We also thank the anonymous referee for helpful comments. The authors were partially supported by the OU Bridge Funding Investment Program.

2 Definitions

2.1 The (pure) surface Houghton group

Here we lay out the definitions necessary for the group $P\mathcal{H}_n$, the pure version of the surface Houghton group defined in [2]. Let $\mathcal{O} = \mathcal{O}_n$ be a sphere with n boundary components, and let T be a torus with two boundary components, ∂^- and ∂^+. Fix an orientation-reversing homeomorphism $\lambda\colon \partial^- \to \partial^+$, and for each i an orientation-reversing homeomorphism μ_i from ∂^- to the i^{th} boundary component of $\mathcal{O}$. We construct Σ_n as follows: begin with $M^1 = \mathcal{O}$, then glue a copy of T to each boundary component of M^1 via the μ_i to obtain M^2. For each $j \geq 2$, glue a copy of T to each boundary component of M^j via λ to obtain M^{j+1}. Then the surface Σ_n is the union of all the M^j's; Σ_n is the surface with n ends, all accumulated by genus. We call $\mathcal{O}$ the *center* of Σ_n, each of the closures of the components of $M^j \setminus M^{j-1}$ is called a *piece*, and each piece B has a canonical homeomorphism $\iota_B\colon B \to T$. We will occasionally write B_k^j to denote the j^{th} piece in the k^{th} end.

Call a subsurface of Σ_n *suited* if it is connected and the union of $\mathcal{O}$ and finitely many pieces. Let $\varphi\colon \Sigma_n \to \Sigma_n$ be a homeomorphism. Call φ *asymptotically rigid* if there exists a suited subsurface $Z \subset \Sigma_n$ (called a *defining surface* for φ) such that

- $\varphi(Z)$ is also suited, and
- φ is *rigid away from* Z, that is, for every piece $B \subset \overline{\Sigma_n \setminus Z}$, we have that $\varphi(B)$ is a piece, and $\varphi|_B = \iota_{\varphi(B)}^{-1} \circ \iota_B$.

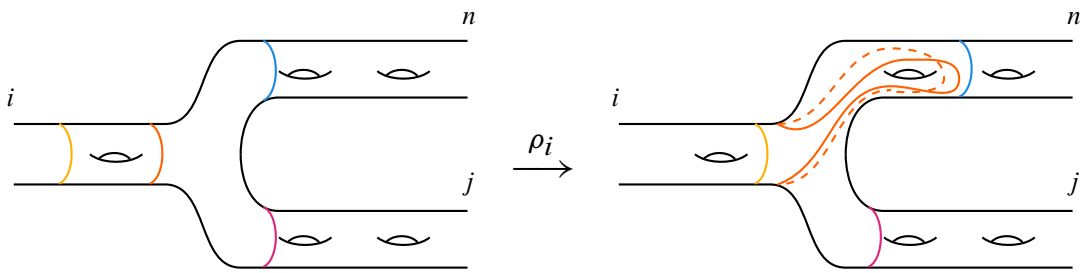

Figure 1: The handle shift ρ_i (which pushes from end i to end n) being applied to various curves in ends i, n, and j.

We may sometimes say that φ is rigid away from a piece B adjacent to $\mathcal{O}$, which we take to mean that φ is rigid away from $\mathcal{O} \cup B$. A special family of suited subsurfaces are those with pieces in only one end. We denote by L_g the suited subsurface consisting of $\mathcal{O}$ and the first g-many pieces in the n^{th} end.

Definition 2.1 (surface Houghton group) The surface Houghton group $\mathcal{H}_n$ is the subgroup of the mapping class group $\mathrm{Map}(\Sigma_n)$ whose elements have an asymptotically rigid representative. The pure surface Houghton group $P\mathcal{H}_n$ is the intersection of $\mathcal{H}_n$ with the pure mapping class group $\mathrm{PMap}(\Sigma_n)$.

We define some special elements of $\mathcal{H}_n$. For σ a permutation of $\{1, \dots, n\}$, choose some homeomorphism of $\mathcal{O}$ permuting the boundary components in the same fashion. For the sake of definiteness, consider σ first as an element of $\mathrm{Map}(S_{0,n})$, the sphere with n punctures; then σ yields a homeomorphism of $\mathcal{O}$, which we also denote by σ. Up to some isotopy, we can assume that the restriction of σ to the i^{th} boundary component of $\mathcal{O}$ is $\mu_{\sigma(i)} \circ \mu_i^{-1}$. Thus, we can obtain a homeomorphism of Σ_n by extending this map to be rigid outside of $\mathcal{O}$; call this extension again by the same name, σ. Such a map is not unique: there are many choices of mapping class in the sphere with n punctures which induce the same permutation.

Next, we define the handle shifts ρ_i, à la [17]. For $i \in \{1, \dots, n-1\}$, we want ρ_i to be a homeomorphism of Σ_n which shifts the i^{th} end towards the n^{th} end by a single genus. Specifically, for $j \geq 1$ map each piece B_n^j to B_n^{j+1}, and map each piece B_i^{j+1} to B_i^j; choose these to be the rigid maps. The ends other than i and n are unchanged, so all that remains is to define how ρ_i acts on $B_i^1 \cup \mathcal{O}$. This is demonstrated by Figure 1. Note that ρ_i is asymptotically rigid outside of $\mathcal{O} \cup B_i^1$.

Lemma 2.2 *Any map $\varphi \in \mathcal{H}_n$ can be written as $\alpha \rho_{i_1} \cdots \rho_{i_j} \sigma$, where α is a compactly supported mapping class, and σ and the ρ_{i_k}'s are as above.*

Proof Consider an arbitrary element $\varphi \in \mathcal{H}_n$: φ clearly induces some permutation σ on the ends of Σ_n; obtain from this $\sigma \in \mathcal{H}_n$. Now we have that $\varphi \circ \sigma^{-1} \in P\mathcal{H}_n$, so we wish to analyze how pure mapping classes act within ends. A rigorous examination of this is contained in Section 2.3. In the meantime, consider how φ acts in the complement of a defining surface: as it does not permute the ends, and takes

pieces to pieces, it must shift each end in or out by some integer amount. From this, we obtain a sequence of handle shifts $\rho_{i_1}, \dots, \rho_{i_k}$. Finally, we have that $\varphi \circ \sigma^{-1} \circ (\rho_{i_1} \cdots \rho_{i_k})^{-1}$ is a compactly supported mapping class, α. □

Remark 2.3 The above expression obtained for φ is far from unique: choices were made both in the selection of σ, as well as in the selection of the handle shifts. However, any different choices made for these will simply change α, and the properties of α are rarely relevant for this paper.

2.2 The contractible cube complex

The group $\mathcal{H}_n$ acts on a contractible cube complex X_n, called the Stein–Farley complex. Consider ordered pairs (Z, φ), where Z is a suited subsurface, and $\varphi \in \mathcal{H}_n$. Declare two such pairs (Z_1, φ_1) and (Z_2, φ_2) to be equivalent if the *transition map* $\varphi_2^{-1} \circ \varphi_1$ takes Z_1 to Z_2 homeomorphically, and is rigid elsewhere (ie in the complement of Z_1).

Denote the equivalence class of (Z, φ) by $[Z, \varphi]$, and the set of equivalence classes by $\mathcal{S}$. The group $\mathcal{H}_n$ acts on $\mathcal{S}$ by

$$\psi \cdot [Z, \varphi] = [Z, \psi \circ \varphi].$$

Define the *complexity* of a pair (Z, φ) to be the genus of Z. This extends to be an $\mathcal{H}_n$-invariant height function $f : \mathcal{S} \to \mathbb{Z} \subset \mathbb{R}$. Given vertices $x_1, x_2 \in \mathcal{S}$, we say that $x_1 \prec x_2$ if there are representatives (Z_i, φ_i) of x_i so that $\varphi_1 = \varphi_2$, $Z_1 \subset Z_2$, and $\overline{Z_2 \setminus Z_1}$ is a (nonempty) *disjoint* union of pieces (ie has at most one piece from each end).

We now construct X_n as a cube complex with vertex set $\mathcal{S}$. We declare that whenever $x_1 \prec x_2$, with $d = f(x_2) - f(x_1)$, there is a d-cube with vertex set given by $\{x \mid x_1 \preceq x \preceq x_2\}$. As Σ_n has n ends, the complex X_n is n-dimensional. The height function f extends affinely to an $\mathcal{H}_n$-invariant height function on X_n. Also, let $X_n^{f \leq k}$ denote the subcomplex spanned by vertices with height at most k. We have the following properties (see [2, Theorem 4.1]):

Theorem 2.4 *The cube complex X_n is contractible, and the action of $\mathcal{H}_n$ satisfies:*

- *The $\mathcal{H}_n$-stabilizers of cubes are finite extensions of mapping class groups of compact surfaces.*
- *For $k \geq 1$, the subcomplex $X_n^{f \leq k}$ is $\mathcal{H}_n$-cocompact.*

We are actually more concerned with the action of $P\mathcal{H}_n$ on X_n. Note that $P\mathcal{H}_n$ is a finite-index subgroup of $\mathcal{H}_n$, and that the finite extension in the above theorem is by permutations of the ends. As mapping class groups of compact surfaces are of type F_∞, we have the following.

Corollary 2.5 *The action of $P\mathcal{H}_n$ on X_n satisfies:*

- *The $P\mathcal{H}_n$-stabilizers of cubes are type F_∞.*
- *For $k \geq 1$, the subcomplex $X_n^{f \leq k}$ is $P\mathcal{H}_n$-cocompact.*

We end this section by demonstrating that vertices and edges can be given a nice form. First, we show that every vertex (indeed, every edge) can be expressed with only elements of $P\mathcal{H}_n$, then we provide a canonical form for the suited subsurface component of the vertex. Specifically, we show that any edge can be represented using only elements in the pure group:

Lemma 2.6 *Any edge in X_n can be expressed as $[Z,\varphi]$—$[Z\cup B,\varphi]$, where $\varphi\in P\mathcal{H}_n$.*

Proof Consider a vertex $[Z,\varphi]$, and let B be a piece adjacent to Z, so that $Z\cup B$ is suited. Write $\varphi=\alpha\rho_{i_1}\cdots\rho_{i_j}\sigma$ as in Lemma 2.2. Then we wish to show that $[Z,\varphi]=[\sigma(Z),\varphi\circ\sigma^{-1}]$, and that $[Z\cup B,\varphi]=[\sigma(Z)\cup\sigma(B),\varphi\circ\sigma^{-1}]$. Consider the diagram

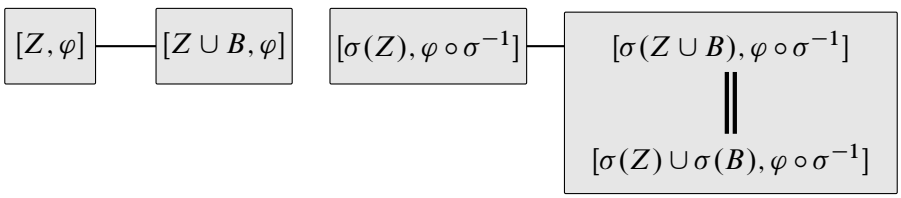

As σ is rigid away from the center $\mathcal{O}$, we have that $\sigma(B)$ is a piece, and of course it must be adjacent to $\sigma(Z)$, so all that remains to be checked is that $[Z,\varphi]=[\sigma(Z),\varphi\circ\sigma^{-1}]$. Note that $(\varphi\circ\sigma^{-1})^{-1}\circ\varphi=\sigma$, which is rigid away from $\mathcal{O}$, and thus takes Z to $\sigma(Z)$ (a suited subsurface), and is certainly rigid outside of Z. □

We assume from here on that any map φ is in $P\mathcal{H}_n$. Vertices have two components, a map and a surface. We now have some control over the map, but what of the surface? For many arguments, it will be convenient to have the following canonical form. Recall that L_g is the surface built by taking the center, and adding in the first g-many pieces in the n^{th} end. Then any edge can be represented with both surfaces an appropriate L_g:

Lemma 2.7 *Any edge in X_n can be expressed as $[L_g,\varphi]$—$[L_{g+1},\psi]$.*

Proof Suppose we are given a vertex $[Z,\varphi']$. Recall that the handle shift ρ_i is rigid away from B_i^1, the piece that it pushes across the center. As $(\varphi)^{-1}\circ\varphi'$ can be chosen to be a composition of handle shifts taking Z to L_g, the only places where $(\varphi)^{-1}\circ\varphi'$ can be nonrigid are pieces in Z. This shows that vertices can be represented as $[L_g,\varphi]$. To extend this to the edge $[Z,\varphi']$ —$[Z\cup B,\varphi']$, consider the diagram

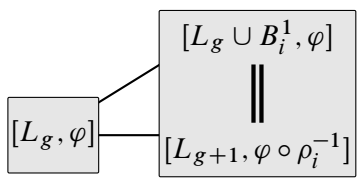

Observe that to obtain an L_{g+1} form for $[Z\cup B,\varphi']$, we can first push Z to L_g, which sends B to a piece B_i^1, and then append one more ρ_i. □

2.3 Characters and Σ-invariants

Recall that a group is of *type* F_m if it admits a proper cocompact action on an $(m-1)$-connected CW-complex. Assume Y is a CW complex and $h\colon Y \to \mathbb{R}$ is continuous. We call the corresponding filtration $(Y^{t\le h})_{t\in\mathbb{R}}$ on Y *essentially* $(m-1)$*-connected* if for any $t \in \mathbb{R}$, there exists $s \le t$ such that inclusion $Y^{t\le h} \hookrightarrow Y^{s\le h}$ induces the trivial map in π_k for $k \le m-1$.

For any group G, we define a *character* of G to be a homomorphism from G to the additive group $\mathbb{R}$. As characters factor through G^{ab} (the abelianization of G), the space of characters is a vector space with the same dimension as the rank of G^{ab}. Excluding the trivial character and modding out by positive scaling, we obtain the *character sphere* $\Sigma(G)$. The Σ-invariants (also called BNS or BNSR invariants, for Bieri–Neumann–Strebel(–Renz)) are a filtration of the character sphere into subspaces

$$\Sigma^0(G) \supseteq \Sigma^1(G) \supseteq \Sigma^2(G) \supseteq \cdots,$$

defined as follows (we use a slight correction of the definition in [19], as suggested in [20]).

Definition 2.8 Let G be a group of type F_m, and let Y be an $(m-1)$-connected CW complex on which G acts cocompactly. Suppose that the stabilizer of any k-cell is of type F_{m-k} and is contained in the kernel of every character of G.[1] For each nontrivial $\chi \in \mathrm{Hom}(G,\mathbb{R})$, there is a character height function $h_\chi\colon Y \to \mathbb{R}$, a continuous map satisfying $h_\chi(gy) = \chi(g) + h_\chi(y)$ for all $y \in Y$ and $g \in G$. Then $\Sigma^m(G)$ is the set of those $[\chi]$ such that the filtration $(Y^{t\le h_\chi})_{t\in\mathbb{R}}$ is essentially $(m-1)$-connected.

Of what use are these invariants, one might ask? A standard result is the following classification of when finiteness properties of G are preserved to subgroups containing the commutator subgroup.

Proposition 2.9 *Let G be a group of type F_m, and let K be a subgroup so that $[G,G] \le K \le G$. Then K is of type F_m if and only if for every character $\chi \in \mathrm{Hom}(G,\mathbb{R})$ such that $\chi(K) = 0$, we have $[\chi] \in \Sigma^m(G)$.*

Definition 2.10 A (*Zaremsky–*)*Morse function* on an affine cell complex Y is a map $h = (\chi, f)\colon Y \to \mathbb{R}\times\mathbb{R}$ such that both χ and f are affine on cells. The codomain is ordered lexicographically, and we require that f take only finitely many values on $Y^{(0)}$, and that there is some $\varepsilon > 0$ such that adjacent vertices v and w satisfy either $|\chi(v) - \chi(w)| \ge \varepsilon$, or $\chi(v) = \chi(w)$ and $f(v) \ne f(w)$.

We record the following general form of the Morse lemma, as well as a more specific version of it. Both will be used in Section 4.

Lemma 2.11 *Let $-\infty \le p \le q \le r \le +\infty$. If for every vertex $v \in Y^{q<\chi\le r}$ the descending link $\mathrm{lk}^{h\downarrow}_{Y^{p\le\chi\le r}}(v)$ is $(k-1)$-connected, then the pair $(Y^{p\le\chi\le r}, Y^{p\le\chi\le q})$ is k-connected. If for every vertex $v \in Y^{p\le\chi<q}$ the ascending link $\mathrm{lk}^{h\uparrow}_{Y^{\chi\le r}}(v)$ is $(k-1)$-connected, then the pair $(Y^{p\le\chi\le r}, Y^{q\le\chi\le r})$ is k-connected.*

[1] That the cell-stabilizers are contained in the kernels is the modification.

Corollary 2.12 *Let $h = (\chi, f)\colon Y \to \mathbb{R}\times\mathbb{R}$ be a Morse function. If Y is $(m-1)$-connected and for every vertex $v \in Y^{\chi<q}$ the ascending link $\mathrm{lk}_Y^{h\uparrow}(v)$ is $(m-1)$-connected, then $Y^{q\leq\chi}$ is $(m-1)$-connected.*

Informally, we can think of the characters on $P\mathcal{H}_n$ as being generated by counting the handle shifts $\rho_1,\dots,\rho_{n-1}$, where ρ_i shifts the i^{th} end into the n^{th} end by one piece. Specifically, for $1\leq i\leq n-1$, take $\chi_i(\varphi)$ to be the negative of the sum of the powers of ρ_i appearing in φ; for $i=n$, take χ_n to be the sum of the powers of ρ_1 through ρ_{n-1}. We stress that this definition, while convenient to use, is not obviously well-defined. An equivalent definition, which is better suited to demonstrating well-definedness, is as follows (details can be found in Section 3 of [5]).

Let γ be an oriented curve that separates one end E of Σ_n from the rest, oriented so that the end E is on the right-hand side of γ. Then γ defines a nonzero element of $H_1^{\text{sep}}(\Sigma_n,\mathbb{Z})$. To every $\varphi\in\mathrm{PMap}(\Sigma_n)$ and $\gamma\in H_1^{\text{sep}}(\Sigma_n,\mathbb{Z})$, associate an integer $\theta_{[\gamma]}(\varphi)$, as a "signed genus" between γ and $\varphi(\gamma)$. Then the map $\theta_{[\gamma]}\colon\mathrm{PMap}(\Sigma_n)\to\mathbb{Z}$ is a well-defined nontrivial homomorphism, depending only on the homology class of γ. By identifying H^1_{sep} with $\mathrm{Hom}(H_1^{\text{sep}},\mathbb{Z})$ via the universal coefficients theorem, we obtain a map $\Theta\colon\mathrm{PMap}(\Sigma_n)\to H^1_{\text{sep}}(\Sigma_n,\mathbb{Z})$, by the rule $\Theta(\varphi)[\gamma]=\theta_{[\gamma]}(\varphi)$. Restricting to $P\mathcal{H}_n$, we obtain characters, and it is not difficult to see that they agree with the informal definition above. We emphasize: χ_i measures how much φ pushes the i^{th} end *out*.

Given a vertex, we can define the character height function $h_{\chi_i}([Z,\varphi])$ by adding the number of pieces of Z in the i^{th} end to $\chi_i(\varphi)$. This is well defined: consider two representatives $[Z,\varphi]=[W,\psi]$. We have that $\varphi^{-1}\psi$ takes W to Z, and is rigid elsewhere. This composition can be further decomposed as a mapping class of some suited subsurface composed with a sequence of handle shifts. This compactly supported mapping class affects neither the distribution of the pieces of Z nor the value of any character on φ, so we need consider only handle shifts. Consider $[Z,\varphi]=[\rho_i(Z),\varphi\circ\rho_i^{-1}]$: we have moved one piece from the i^{th} end to the n^{th} end, while increasing the amount of pushing into the i^{th} end and out of the n^{th} end by one. These cancel out in h_{χ_i} and h_{χ_n}. All other basis characters are unchanged under these operations, and thus we have well-defined character height functions on X_n. For an arbitrary character χ, we shall henceforth abuse notation by writing χ when we mean h_χ. The domain will typically be clear.

3 The Stein–Farley complex X_n is CAT(0)

We have a contractible cube complex on which $P\mathcal{H}_n$ acts nicely. In this section, we show that X_n is in fact CAT(0). To begin, we require the following definition.

Definition 3.1 A cube complex X is *cube-complete* if whenever $X^{(1)}$ contains an embedded copy of the 1-skeleton of a d-cube, X contains the entire d-cube.

Once we have shown that X_n is cube-complete, we seek to apply the following proposition, which is a version of [14, Proposition 4.6], where it is taken out of the proof of [9, Theorem 6.1].

Proposition 3.2 *Let X be a cube-complete cube complex, whose 1-skeleton $X^{(1)}$ is a graph with no loops or bigons. Suppose that*

(a) *X is simply connected,*

(b) *$X^{(1)}$ satisfies the 3-square condition.*

Then X is CAT(0).

The 3-square condition says that whenever the cube complex has 3 squares intersecting in a vertex, and pairwise intersecting in edges (call such an arrangement a 3-wheel), then they are part of a full 3-cube.

Remark 3.3 In [14], there is a third condition: that the 1-skeleton not contain a copy of the complete bipartite graph $K_{2,3}$. This third condition is redundant: A careful examination of the proof of Theorem 6.1 in [9], especially the paragraph beginning "To prove the converse", yields that the other hypotheses imply the nonexistence of $K_{2,3}$'s: A $K_{2,3}$ is a 3-wheel where all the vertices opposite the common vertex are identified, which cannot be the case in an embedded 3-cube. Another, smaller modification: there was a condition that the 1-skeleton be a connected graph, which is automatically true in a simply connected cube complex.

As we have a height function for X_n, and as adjacent vertices always differ by exactly 1, there are two ways in which X_n can fail to be cube-complete.

Definition 3.4 By a *collapsed* cube, we mean the 1-skeleton of a d-cube, such that the difference between its highest and lowest heights is less than d. By an *empty* cube, we mean the 1-skeleton of a d-cube that is not filled by a d-cube.

Given the definition of X_n, we need to check two things: that $X_n^{(1)}$ has no collapsed cubes, and that $X_n^{(1)}$ has no empty cubes.

Choosing vertex representatives to have their homeomorphism components contained in the pure group allows us a consistent definition of "direction" for adding and removing pieces, ie we may consistently label the ends from 1 to n without worrying about any potential reindexing.

In order to show that X_n is cube-complete, we shall proceed as follows: first, we show that given a vertex $[Z,\varphi]$, there is only one ascending edge per direction (see Lemma 3.5). Secondly, we show that there are no collapsed squares. This easily gives that there are no collapsed cubes of any dimension. Finally, we show that there are no empty squares, and from this induct to show that there are no empty cubes.

Lemma 3.5 (uniqueness of ascending edges) *Given a vertex $v=[Z,\varphi]$, there exists exactly one ascending edge in each direction, ie for $i\in\{1,\dots,n\}$, the collection $\{[Z\cup B_i,\varphi]\}$, where B_i is the piece in the i^{th} end adjacent to Z, exhausts the vertices of the ascending link of v.*

Proof What we wish to show is that, given two representatives $[Z,\varphi]=[Z',\varphi']$, whenever we ascend from $[Z',\varphi']$ by adding a piece in the i^{th} end, we obtain the same vertex as by adding a piece in the i^{th} end to $[Z,\varphi]$. Consider the diagram

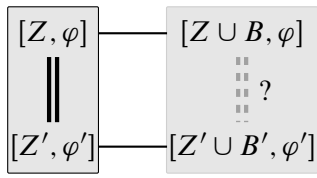

Our goal is to show that, assuming B' and B are both in the i^{th} end, we have $[Z'\cup B',\varphi']=[Z\cup B,\varphi]$. As $(\varphi')^{-1}\circ\varphi$ is rigid outside Z, it takes B to some piece adjacent to Z'. As $(\varphi')^{-1}\circ\varphi\in P\mathcal{H}_n$, this implies that $(\varphi')^{-1}\circ\varphi(B)$ must still be in the i^{th} end (elements of $P\mathcal{H}_n$ cannot rigidly move pieces from one end into another), and hence is B'. We then have that $(\varphi')^{-1}\circ\varphi$ takes $Z\cup B$ to $Z'\cup B'$, and is rigid elsewhere, ie that $[Z\cup B,\varphi]=[Z'\cup B',\varphi']$. Thus, there can be only one ascending edge per direction. □

Combining this with the L_g representatives of Lemma 2.7, we see that the ascending edges over $[L_g,\varphi]$ are of the form $[L_{g+1},\varphi\rho_i^{-1}]$, for $i\in\{1,\dots,n\}$. (Recall that ρ_n is just the identity.) Next, we prove that there are no collapsed squares.

Lemma 3.6 *Given any square in $X_n^{(1)}$, the difference between the maximal and minimal heights of vertices is 2.*

Proof For the sake of contradiction, suppose there is some collapsed square, which has only two values for the heights of its vertices. Writing the lower height vertices as $[L_g,\varphi]$ and $[L_g,\psi]$, we see that a collapsed square must occur as in Figure 2. That the same representative can be used for both edges from one of the side vertices is justified by Lemma 3.5. To show that the collapsed square in Figure 2 cannot occur, we assume that the middle pair of vertices are distinct, and prove that the leftmost and rightmost vertices cannot be distinct. The equality of the different forms of the middle vertices says that the transition map $\rho_{i_k}\varphi^{-1}\psi\rho_{j_k}^{-1}$ takes L_{g+1} to itself and is rigid elsewhere, for $k=1,2$. This implies that $\rho_{i_k}\varphi^{-1}\psi\rho_{j_k}^{-1}$ has net-zero shifting in all ends, and hence is compactly supported (recall that for asymptotically rigid maps, these are the same). As the middle vertices are distinct, we know that $i_1\neq i_2$ and $j_1\neq j_2$. Thus, we see that $i_k=j_k$ for $k=1,2$.

We will now show that $[L_g,\varphi]=[L_g,\psi]$ by considering the transition map $\varphi^{-1}\psi$ on pieces outside of L_g. The rigidity of the $\rho_{i_k}\varphi^{-1}\psi\rho_{j_k}^{-1}$ maps outside L_{g+1} immediately handles all but three pieces: $B_{i_1}^1$, $B_{i_2}^1$, and B_n^{g+1}. The last of these is again easy: the L_{g+1} transition maps are rigid on the $(g+2)^{\text{nd}}$ piece, which they first push to the $(g+1)^{\text{st}}$ piece, then apply $\varphi^{-1}\psi$, then push back out by one. As this must be sent rigidly to the $(g+2)^{\text{nd}}$ piece, we see that $\varphi^{-1}\psi$ must be rigid on the $(g+1)^{\text{st}}$ piece. For $B_{i_1}^1$, consider that $\varphi^{-1}\psi$ and $\rho_{i_2}\varphi^{-1}\psi\rho_{i_2}^{-1}$ act the same on it. As the latter is rigid there, so is the former. Symmetrically, we have rigidity on $B_{i_2}^1$, and are done. Specifically, we have shown that $\varphi^{-1}\psi$ takes L_g to itself, and is rigid elsewhere, implying that the supposed collapsed square was degenerate to begin with. □

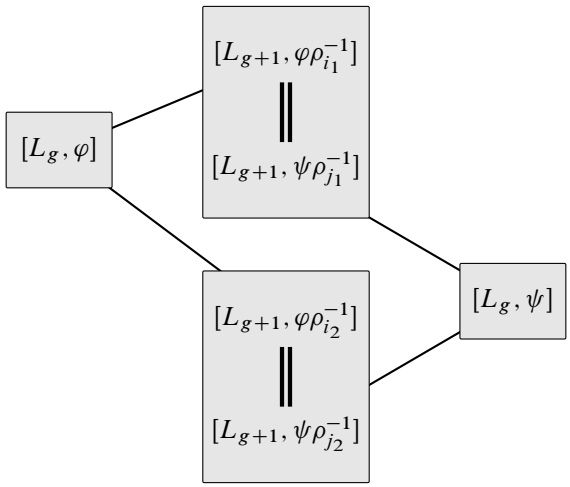

Figure 2: A collapsed square.

Proposition 3.7 *The Stein–Farley complex X_n is cube complete.*

Proof We begin by observing that there can be no collapsed cubes: any such cube must contain a collapsed square. Now we show that there are no empty squares. As there are no collapsed squares, we know that the 1-skeleton of a square must have a lowest vertex, say at height g, two vertices at height $g+1$, and one vertex at height $g+2$. By the uniqueness (per end) of ascending edges (Lemma 3.5), we see that given a representative $[Z,\varphi]$ for the height g vertex, the other vertices must be of the form $[-,\varphi]$, with the blank filled by $Z\cup B_1$, $Z\cup B_2$, and $Z\cup B_1\cup B_2$, depending on height. The only concern about the $g+2$ vertex would be that the pieces lie in the same end, but we see that this cannot occur. Thus, there are no empty squares.

Finally, we show that there are no empty cubes of any dimension. We do so inductively: suppose we already have that there are no empty $(d-1)$-cubes. Consider the 1-skeleton of a d-cube C. Let the bottom vertex of C be $[Z,\varphi]$, at height g. Each of the vertices of C adjacent to $[Z,\varphi]$ is of the form $[Z\cup B_i,\varphi]$, where B_i is in the i^{th} end; each i appears at most once by Lemma 3.5. Also, by the inductive hypothesis, we have that $[Z,\varphi]$ connects to each of the vertices at height $g+d-1$ via a $(d-1)$-cube, so that these vertices can be written as $[Z\cup B_{i_1}\cup\cdots B_{i_{j-1}}\cup\widehat{B}_{i_j}\cup B_{i_{j+1}}\cup\cdots\cup B_{i_d},\varphi]$, where the hat indicates omission. Choose two such vertices, and consider their common lower vertex (the one missing the two pieces missing in either of the chosen vertices). Applying the doctrine of no empty squares to these 3 vertices and the apex yields that the apex can be written as $[Z\cup B_{i_1}\cup\cdots\cup B_{i_d},\varphi]$, and thus the cube is filled. □

Remark 3.8 The argument above that ascending edges are unique per direction emphatically does not hold for descending edges. In fact, for any vertex of height at least 1, there are infinitely many descending edges per direction in which it can have a descending edge (given $[Z,\varphi]$, choose $\psi\in\mathcal{PH}_n$ such that $\varphi^{-1}\psi|_Z$ is a homeomorphism, with some nontrivial behavior in the piece to be removed, and such that

$\varphi^{-1}\psi$ is rigid outside Z). It is helpful to keep in mind the following: when ascending, one can always choose coherent representatives from a minimal vertex; when descending, nothing is guaranteed.

Theorem 3.9 *Let X_n denote the Stein–Farley complex for $\mathcal{H}_n$. Then X_n is* CAT(0).

Proof By Proposition 3.7, X_n is cube-complete. By construction, $X_n^{(1)}$ has no loops or bigons. As X_n is contractible, it is simply connected, so item (a) is covered. For item (b), we need to show that 3-wheels can always be completed. There are four possibilities: the common vertex can be lowest height, highest height, or either of the intermediate heights. In the first case, we can have all three edges out of the common vertex use the same representative, each adding a piece/shift in a different end, so that the cube is completed by adding in all three at once. In the intermediate height cases, something similar occurs: we always have the lowest height vertex of the desired cube from which we can build upwards. It is useful to keep in mind the following arrangement of the 1-skeleton of a 3-cube; the solid lines indicate the 3-wheel, and the dashed lines indicate the rest of the 3-cube:

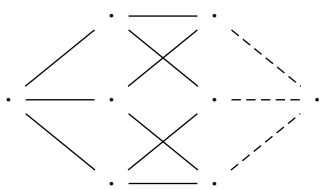

The difficulty is in the remaining case, where the common vertex is maximal in the cube. We begin with the diagram in Figure 3. Our goal is to obtain a representative $[L_{g+2}, \varphi \circ \rho_k]$ for the height $g+2$ vertex currently represented with both φ' and φ''. This will imply the existence of a vertex $[L_g, \varphi \circ \rho_i \rho_j \rho_k]$, which realizes this 3-wheel in a 3-cube. To begin, we have the following lemma that justifies part of the diagram.

Lemma 3.10 *Whenever $[L_g, \varphi \circ \rho_i] = [L_g, \varphi' \circ \rho_j]$ and $[L_{g+1}, \varphi] = [L_{g+1}, \varphi']$, we have that $i = j$.*

Proof First, observe that $[L_g, \varphi \circ \rho_i] = [L_g, \varphi' \circ \rho_j]$ means that $h = \rho_i^{-1}\varphi^{-1}\varphi'\rho_j$ takes L_g to itself, and is rigid elsewhere; also, $[L_{g+1}, \varphi] = [L_{g+1}, \varphi']$ means that $\varphi^{-1}\varphi'(L_{g+1}) = L_{g+1}$. Assume $j \neq n$, and consider the piece B_j^1; let α be the boundary curve it shares with $\mathcal{O}$. What does h do to α? First, ρ_j deforms it into the center and stretches it into the n^{th} end, making it an essential curve in L_g. Then $\varphi^{-1}\varphi'$ moves it around to some essential curve in L_{g+1}. Finally, ρ_i^{-1} pushes and stretches it towards the i^{th} end. As B_j^1 is not in L_g, it must be sent to a piece, and $h(\alpha)$ must be the intersection of this piece with L_g. As all representatives here do not permute ends, it must be that $h(\alpha) = \alpha$. If $i \neq j$, then ρ_i fixes α, so $\rho_i(\alpha) \neq \varphi^{-1}\varphi'\rho_j(\alpha)$, and $h(\alpha) \neq \alpha$. Hence we must have that $i = j$. In the case where $j = n$, equality is obvious. □

Now, we seek the following equality: $[L_{g+2}, \varphi \circ \rho_k] = [L_{g+2}, \varphi' \circ \rho_k]$. That is, we wish to show that $\rho_k^{-1}\varphi^{-1}\varphi'\rho_k$ takes L_{g+2} to itself, and is rigid elsewhere. We already know that this holds for $\rho_i^{-1}\varphi^{-1}\varphi'\rho_i$. The rigidity here occurs precisely when $\varphi^{-1}\varphi'$ does nothing untoward to whatever single piece gets

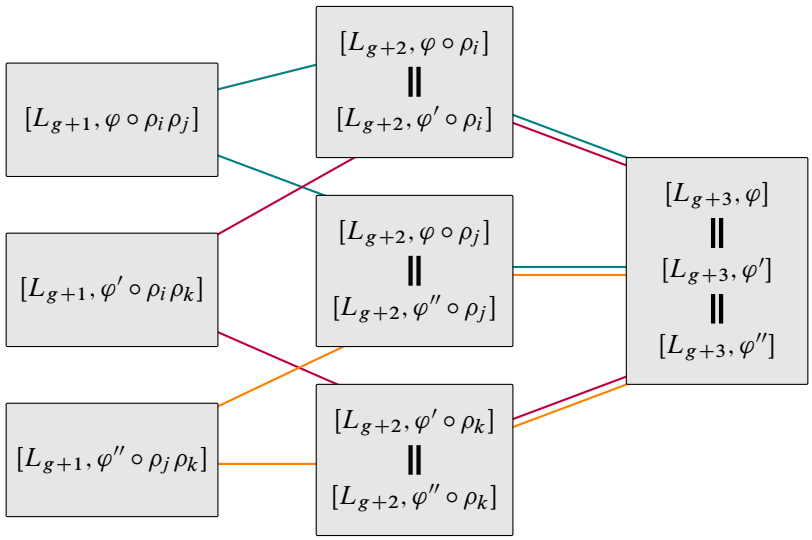

Figure 3: A 3-wheel with common vertex maximal.

pushed into the n^{th} end, so it doesn't matter which index we choose! (If we have chosen the index n, there is no pushing. To resolve this, simply choose a different height $g+2$ vertex to work with. They can't share ends, as they are in squares with lower vertices, so we can apply Lemma 3.5.)

One small detail remains: we have a bottom vertex connecting to one of the height $g+1$ vertices, but does it connect to the other height $g+1$ vertices? Suppose not: then the cube between $[L_{g+3}, \varphi]$ and $[L_g, \varphi \circ \rho_i \rho_j \rho_k]$ contains a square differing from one of the original 3 squares only at its height $g+1$ vertex (that is, at its lowest vertex). This however means that we would have two squares agreeing on 3 vertices, disagreeing on the 4th, which would contain a collapsed square. As this cannot occur (Lemma 3.6), our height g vertex connects to each of the height $g+1$ vertices we started with. □

4 Σ-invariants of $P\mathcal{H}_n$

The methods of this section are taken from [19; 20]. They need only minimal modification to work in this setting, and are presented partly for the sake of having the entire argument in one place. We begin by noting that the abelianization of $P\mathcal{H}_n$ is $\mathbb{Z}^{n-1}$, and the abelianization map is $(\chi_1, \dots, \chi_{n-1})$ (see Section 6 of [2]). As $\chi_1 + \cdots + \chi_n = 0$, any (nontrivial) character χ can be written (up to renumbering the ends of Σ_n) in ascending standard form, ie $\chi = a_1\chi_1 + \cdots + a_n\chi_n$, with $a_1 \le \cdots \le a_{m(\chi)} < a_{m(\chi)+1} = \cdots = a_n = 0$. We shall henceforth assume all characters are written in such form. Observe that $m(\chi)$ is defined to be the maximal index so that $a_{m(\chi)} < a_n = 0$. Then our goal in this section is to prove the following.

Theorem 4.1 *Let χ be a nonzero character of $P\mathcal{H}_n$. Then $[\chi] \in \Sigma^{m(\chi)-1}(P\mathcal{H}_n) \setminus \Sigma^{m(\chi)}(P\mathcal{H}_n)$.*

4.1 Inclusion

We begin by showing the inclusion into the $(m(\chi)-1)$-layer of $\Sigma(P\mathcal{H}_n)$.

Theorem 4.2 *For any nonzero character $\chi \in \mathrm{Hom}(P\mathcal{H}_n, \mathbb{R})$, we have $[\chi] \in \Sigma^{m(\chi)-1}(P\mathcal{H}_n)$.*

We recall here the form of the Morse lemma we shall use in this section.

Corollary 4.3 *Let $h = (\chi, f): Y \to \mathbb{R}\times\mathbb{R}$ be a Morse function. If Y is $(m-1)$-connected and for every vertex $v \in Y^{\chi<q}$ the ascending link $\mathrm{lk}_Y^{h\uparrow}(v)$ is $(m-1)$-connected, then $Y^{q\le\chi}$ is $(m-1)$-connected.*

From now on, we write X for X_n. The role of Y shall be played by sublevel subcomplexes $X^{f\le q}$ for sufficiently large q. Let $\chi = a_1\chi_1 + \cdots + a_n\chi_n$ be a nontrivial character of $P\mathcal{H}_n$, written in ascending standard form. The function $h = (\chi, f): Y \to \mathbb{R}^2$ is a Morse function. Observe that between any two adjacent vertices, each basis character differs by 1, 0, or -1. We examine h-ascending links in Y.

Lemma 4.4 *Let $v = [L_g, \varphi]$ be a vertex in Y. An adjacent vertex $w = [L_{g\pm1}, \psi]$ is in the h-ascending link of v if and only if one of the two following conditions holds*:

- $w = [L_{g-1}, \psi]$, *with* $v = [L_g, \psi\rho_i^{-1}]$ *where* $i \le m(\chi)$.
- $w = [L_{g+1}, \varphi\rho_i^{-1}]$, *where* $i \ge m(\chi)+1$.

Proof If w is f-ascending from v, then $w = [L_{g+1}, \varphi\rho_i^{-1}]$ for some i. If $i \ge m(\chi)+1$, then $\chi(w) = \chi(v)$, so w is h-ascending from v. If $i \le m(\chi)$, then w is χ-descending from v, hence h-descending from v. If w is f-descending, then $v = [L_g, \psi\rho_i^{-1}]$ for some i. If $i \ge m(\chi)+1$, then we again have $\chi(w) = \chi(v)$, so that w is h-descending from v. If $i \le m(\chi)$, then w is χ-ascending from v, hence h-ascending. □

Remark 4.5 It is worth observing a difference here from the situation in [19]. For the regular Houghton groups, there is one way to go "up" and finitely many ways to go "down" per end (with respect to f); here, while there is only one way to go up per end, there are infinitely many ways to go down. However, the h-ascending link is still a join of its intersection with the f-ascending and f-descending links, as these elements cannot share ends.

We shall need the following fact (see [2, Section 5.2]): Let $v \in X_n$ be a vertex. If $f(v) \ge 2n$, then the f-descending link of v is $(n-2)$-connected. The version of this statement for the regular Houghton groups uses $f(v) \ge 2n-1$ (see [16, Theorem 3.52]); this discrepancy will result in most of the numbers in the following proposition being either one above or one below the corresponding numbers in [19].

Now set $q = 3n-2$, so $Y = X^{f\le 3n-2}$.

Proposition 4.6 *Let v be a vertex in Y. Then $\mathrm{lk}_Y^{h\uparrow}(v)$ is $(m(\chi)-2)$-connected.*

Proof We have that $f(v)$ is between 0 and $3n-2$. Suppose that $f(v) \leq 2n + m(\chi) - 2$. Writing $v = [L_g, \varphi]$, we have that there are $n - m(\chi)$ indices i for which $v' = [L_{g+1}, \varphi \circ \rho_i^{-1}]$ is h-ascending. As $(2n + m(\chi) - 2) + (n - m(\chi)) = 3n - 2$, the entire f-ascending link of v in X is contained in Y. Thus, the f-ascending part of the h-ascending link is an $(n-m(\chi)-1)$-simplex, which is contractible, so that $\mathrm{lk}_Y^{h\uparrow}(v)$ is contractible.

Suppose instead that $2n + m(\chi) - 1 \leq f(v) \leq 3n - 2$, and thus that Y does not contain the entire f-ascending part of the h-link of v. We still have its $(3n-f(v)-3)$-skeleton, which, being a skeleton of an $(n-m(\chi)-1)$-simplex, is $(3n-f(v)-4)$-connected. As $f(v) \geq 2n$, we have that the entire f-descending part of the h-ascending link in is Y. This is isomorphic to the f-descending link of a vertex of the same f-height as v in $X_{m(\chi)}$, which we know is $(m(\chi)-2)$-connected so long as $f(v) \geq 2m(\chi)$. As $f(v) \geq 2n + m(\chi) - 1$, we need that $2n - 1 \geq m(\chi)$, which is certainly true. The join is now $((3n-f(v)-3)+(m(\chi)-1))$-connected. As $f(v) \leq 3n-2$, we have that $(3n-f(v)+m(\chi)-4) \geq m(\chi)-2$, and thus that $\mathrm{lk}_Y^{h\uparrow}(v)$ is $(m(\chi)-2)$-connected. □

4.2 Exclusion

To demonstrate that the result of the last section is sharp, we shall need to employ different techniques. The main result of this section is:

Theorem 4.7 *For any nonzero character $\chi \in \mathrm{Hom}(P\mathcal{H}_n, \mathbb{R})$, we have $[\chi] \notin \Sigma^{m(\chi)}(\mathcal{H}_n)$.*

We gather here the propositions and definitions from [20] that we will need, starting with the strong nerve lemma:

Proposition 4.8 *Let X be a CW-complex covered by subcomplexes $(X_i)_{i\in I}$ and let L be the nerve of the cover. Let $n \geq 1$. Suppose that any nonempty intersection $X_{i_1} \cap \cdots \cap X_{i_r}$ is $(n-r)$-connected. Then $H_k(X) \cong H_k(L)$ for all $k \leq n-1$, and $H_n(X)$ surjects onto $H_n(L)$.*

Definition 4.9 For a set of indices $K \subseteq [n]$, consider the subcomplex $\bigcap_{i\in K} X^{\chi_i \leq 0}$ of X. Call any connected component of such a subcomplex a *K-blanket*. By a *blanket* we mean a K-blanket for some unspecified K.

Recall that a subcomplex Z of a CAT(0) cube complex Y is *locally combinatorially convex* if every link in Z of a vertex $z \in Z^{(0)}$ is a full subcomplex of the link of z in Y, and *combinatorially convex* if it is connected and locally combinatorially convex. It is known that combinatorially convex implies CAT(0), hence contractible. In particular, this applies to connected components of locally combinatorially convex subcomplexes.

Lemma 4.10 *For any K, $\bigcap_{i\in K} X^{\chi_i \leq 0}$ is locally combinatorially convex. Thus, blankets are combinatorially convex, and hence* CAT(0) *and contractible.*

Proof It suffices to show that each $X^{\chi_i \leq 0}$ is locally combinatorially convex. Given a pair of adjacent vertices, we can write them as $v = [L_g, \varphi]$ and $w = [L_{g+1}, \varphi \circ \rho_j^{-1}]$ for some j. Then $\chi_i(w) - \chi_i(v) = \delta_{i,j}$. Thus, if C is a cube containing v, and $w_1, \ldots, w_k$ are the vertices of C adjacent to v, then the maximum and minimum values of χ_i on C lie in $\{\chi_i(v), \chi_i(w_1), \ldots, \chi_i(w_k)\}$. Thus, whenever $v \in X^{\chi_i \leq 0}$ and all these w_j's lie in the link of v in $X^{\chi_i \leq 0}$, then the cube C lies in $X^{\chi_i \leq 0}$. This implies that the link of v in $X^{\chi_i \leq 0}$ is a full subcomplex of the link of v in X. □

As in [20], we have the immediate corollary that intersections of blankets are blankets for the union of the sets of indices. What follows is essentially identical to Zaremsky's approach for the Houghton groups, reproduced here with minor changes for the sake of completeness. Recall now the more general statement of the Morse lemma, Lemma 2.11. For notational convenience, we write $X_{f \leq k}$ for $X^{f \leq k}$, and $X_{f \leq k}^{t \leq \chi}$ for the intersection $X_{f \leq k} \cap X^{t \leq \chi}$.

Lemma 4.11 *If $X_{f \leq 3n-2}^{0 \leq \chi}$ is not $(m(\chi)-1)$-connected, then $[\chi] \in \Sigma^{m(\chi)}(P\mathcal{H}_n)^c$.*

Proof If $[\chi] \in \Sigma^{m(\chi)}(P\mathcal{H}_n)$, then the filtration $(X_{f \leq 3n-2}^{t \leq \chi})_{t \in \mathbb{R}}$ is essentially $(m(\chi)-1)$-connected. Every h-ascending link of a vertex in $X_{f \leq 3n-2}$ is $(m(\chi)-2)$-connected, so for any $s \leq t$ the inclusion $X_{f \leq 3n-2}^{t \leq \chi} \hookrightarrow X_{f \leq 3n-2}^{s \leq \chi}$ induces an isomorphism in π_k for $k \leq m(\chi) - 2$, and a surjection in $\pi_{m(\chi)-1}$. By assumption, for any t there is some $s \leq t$ such that this inclusion induces a trivial map in π_k for $k \leq m(\chi) - 1$, implying that $X_{f \leq 3n-2}^{s \leq \chi}$ is $(m(\chi)-1)$-connected. Rescaling if necessary, we can assume that $s \in \chi(P\mathcal{H}_n)$, and thus we can translate to obtain $X_{f \leq 3n-2}^{s \leq \chi} \cong X_{f \leq 3n-2}^{0 \leq \chi}$, so that $X_{f \leq 3n-2}^{0 \leq \chi}$ is itself $(m(\chi)-1)$-connected. □

In order to show that $X_{f \leq 3n-2}^{0 \leq \chi}$ is not $(m(\chi)-1)$-connected, we will apply the strong nerve lemma to a covering we now define. For $1 \leq i \leq n$, let $\{Z_i^\alpha\}$ be the collection of $\{i\}$-blankets in X. The α's are indices in some set, and this index set is not itself important. Set also

$$Y_i^\alpha = Z_i^\alpha \cap X_{f \leq 3n-2}^{0 \leq \chi}.$$

Lemma 4.12 *The Y_i^α with $1 \leq i \leq m(\chi)$ cover $X_{f \leq 3n-2}^{0 \leq \chi}$.*

Proof As the Z_i^α are the connected components of the $X^{\chi_i \leq 0}$, it suffices to show that

$$X^{0 \leq \chi} \subseteq \bigcup_{i=1}^{m(\chi)} X^{\chi_i \leq 0}.$$

As $\chi = a_1\chi_1 + \cdots + a_{m(\chi)}\chi_{m(\chi)}$, with all coefficients negative, any vertex $v \in X$ with $\chi(v) \geq 0$ must satisfy $\chi_i(v) \leq 0$ for some i. This implies the inclusion on vertices. Given a cube in $X^{0 \leq \chi}$, let v be its maximal vertex with respect to f. Then all the vertices w of the cube satisfy $\chi_i(w) \leq \chi_i(v)$, and hence the entire cube lies in whichever $X^{\chi_i \leq 0}$ contains v. □

Lemma 4.13 *Any nonempty intersection of subcomplexes of the form Y_i^α with $1 \leq i \leq m(\chi)$ is $(m(\chi)-2)$-connected.*

Proof To be nonempty, such an intersection can include at most one term Y_i^α for each i, and can thus be written $Y = Y_{i_1}^{\alpha_1} \cap \cdots \cap Y_{i_r}^{\alpha_r}$, with the i_j all pairwise distinct. Let $Z = Z_{i_1}^{\alpha_1} \cap \cdots \cap Z_{i_r}^{\alpha_r}$, so that $Y = Z \cap X^{0\le\chi}_{f\le 3n-2}$. We apply Morse-theoretic techniques to Z, this time using Lemma 2.11. As Z is an intersection of blankets, it is a blanket, and thus contractible. Given adjacent vertices $w = [L_g, \varphi]$ and $v = [L_{g+1}, \varphi \circ \rho_i^{-1}]$ with $v \in Z$, we have that $w \in Z$. Thus, for any vertex of Z, the entire f-descending link is in Z. As this is $(n-2)$-connected for $f(v) \ge 2n$, we see that $Z_{f\le 3n-2}$ is $(n-2)$-connected, and so is certainly $(m(\chi)-2)$-connected. Now consider Y as $Z^{0\le\chi}_{f\le 3n-2}$. As before, the h-ascending link of a vertex v is a join between its f-ascending and f-descending parts. The latter is in Z for the same reasons as above; the former is in Z because it consists of directions i where $m(\chi)+1 \le i \le n$, on which the considered χ_{i_j} are constant. Thus, the h-ascending link of v is in $Z_{f\le 3n-2}$. As $Z_{f\le 3n-2}$ is $(m(\chi)-2)$-connected, the Morse lemma tells us that Y is $(m(\chi)-2)$-connected. □

Let L be the nerve of the covering of $X^{0\le\chi}_{f\le 3n-2}$ by the Y_i^α. Since $[\chi] \in \Sigma^{m(\chi)-1}(P\mathcal{H}_n)$, we know that $X^{0\le\chi}_{f\le 3n-2}$ is $(m(\chi)-2)$-connected, so by the strong nerve lemma L is $(m(\chi)-2)$-connected. The final missing piece is to prove that L is not $(m(\chi)-1)$-acyclic.

Lemma 4.14 *The nerve L is not $(m(\chi)-1)$-acyclic.*

Proof Consider vertices corresponding to Y_i^α and Y_j^β. These vertices can only be adjacent if $i \ne j$, so L is $(m(\chi)-1)$-dimensional. Thus, it suffices to exhibit a nontrivial $(m(\chi)-1)$-cycle. This will come from a collection of $2m(\chi)$ vertices, 2 for each $i \in \{1, \dots, m(\chi)\}$, labeled as $Y_i^{\epsilon_i}$, with $\epsilon \in \{1,2\}$, with the property $Y_1^{\epsilon_i} \cap \cdots \cap Y_{m(\chi)}^{\epsilon_{m(\chi)}} \ne \varnothing$. This will yield an embedded $(m(\chi)-1)$-sphere in L, which is homologically nontrivial for dimensional reasons.

Recall that $\mathcal{O}$ is the centerpiece of our surface. For each i, take Y_i^1 to be the Y_i^α containing $v_0 = [\mathcal{O}, \mathrm{id}]$, and take Y_i^2 to be the Y_i^α containing $v_i = [\mathcal{O}, \varphi_i]$, where φ_i is some nontrivial mapping class in the i^{th} end; for the sake of definiteness, choose some essential simple closed curve in the standard piece, and let φ_i be the Dehn twist about its image in B_i^1. Any intersection $Y_1^{\epsilon_1} \cap \cdots \cap Y_{m(\chi)}^{\epsilon_{m(\chi)}}$ includes $w = \left[\mathcal{O}, \prod \varphi_i\right]$, where the product (in an arbitrary order) is taken over those i with $\epsilon_i = 2$, and is therefore nonempty.

It remains to show that $Y_i^1 \ne Y_i^2$ for each i. It suffices to show that $Z_i^1 \ne Z_i^2$. If these are equal (call them Z_i), then we can connect v_0 to v_i via a path in Z_i. In X, one could assume that such a path is one on which f first strictly increases, then strictly decreases; since Z_i is combinatorially convex, this property holds for Z_i as well. Since the path lies in Z_i, χ_i is nonpositive on the whole path. By the uniqueness of ascending edges, this is a path of the form

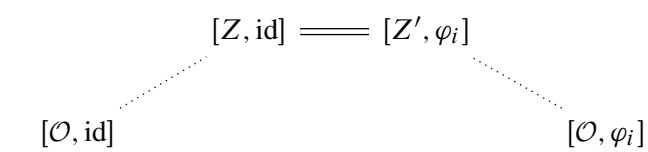

where each of the ascending dotted lines indicates a sequence of edges which only add pieces. Since $\chi_i(v_0) = 0 = \chi_i(v_i)$, none of the edges of the path can be obtained by adding a piece in the i^{th} end. One of Z and Z' must have at least one piece in the i^{th} end, as the transition map is φ_i, which has nonrigid behavior in the i^{th} end. Thus, we have a contradiction. □

5 Subgroups of maximal finiteness length

We use the notation $\mathrm{fl}(G) = n$ to mean that G is a group of type F_n but not of type F_{n+1}. We know that if $G < H$ is finite index, then $\mathrm{fl}(G) = \mathrm{fl}(H)$. But what of the converse, that is, when does $\mathrm{fl}(G) = \mathrm{fl}(H)$ imply that G is finite index in H? For Houghton groups, and for the (pure) surface Houghton groups, the answer is positive for sufficiently large subgroups, in particular for coabelian subgroups. We denote by G' the commutator subgroup of a group G. Recall that the commutator subgroups of the Houghton group H_n and the pure surface Houghton group $P\mathcal{H}_n$ are the finitely supported and compactly supported elements, respectively. (Really, we could say compactly supported in both cases, using the discrete topology for the $\mathbb{N}$-rays on which the Houghton group acts.) For Theorem 5.1 and Proposition 5.2, let H be either H_n or $P\mathcal{H}_n$. For H_n, this is a mild extension of [20, Corollary 2.6].

Theorem 5.1 *Let $G < H$ be a subgroup intersecting the commutator subgroup H' in a finite-index subgroup. If $\mathrm{fl}(G) = n-1$, then G is finite index in H.*

It suffices to show this for groups containing the commutator subgroup. We begin by showing that when G contains the commutator subgroup, the image of G in the abelianization is a maximal rank sublattice, which will then imply finite index.

Proposition 5.2 *Let $G < H$ be as above, ie $\mathrm{fl}(G) = n-1$, and G contains the commutator H'. Write F for the abelianization map to $\mathbb{Z}^{n-1}$. Then $F(G)$ is finite index in $\mathbb{Z}^{n-1}$.*

Proof We start by observing that $F = (\chi_1, \dots, \chi_{n-1})$. As $\Sigma^{n-1}(H)$ is empty, the only way for G to be of type F_{n-1} is that it cannot be killed by any nonzero character of H. We shall construct a sequence of maps $F_i = (F_i^1, \dots, F_i^{n-1}) : H \to \mathbb{Z}^{n-1}$ with $F = F_1$, and a sequence of elements $g_i \in G$ such that $F_n^j(g_i)$ is positive when $i = j$, and zero when $i < j$. Then $F_n(G)$ will be a maximal rank sublattice of $\mathbb{Z}^{n-1}$, obtainable from $F_1(G)$ by integer matrix transformations.

As $\chi_1(G) \neq 0$, we have some smallest positive integer $c_1 \in \chi_1(G)$. Let $g_1 \in \chi_1^{-1}(c_1) \cap G$, and write $F_1(g_1) = (a_1^1, \dots, a_{n-1}^1)$ (note that $a_1^1 = c_1$). Set $F_2 = (\chi_1, a_1^1\chi_2 - a_2^1\chi_1, \dots, a_1^1\chi_{n-1} - a_{n-1}^1\chi_1)$. Then $F_2(g_1) = (c_1, 0, \dots, 0)$. By the same argument, we can choose $c_2 > 0$ minimal in the image of the character $a_1^1\chi_2 - a_2^1\chi_1$, and then $g_2 \in G$ in its preimage. Carrying on, we obtain F_i from F_{i-1} by modifying only the components from i onwards, and build a collection $g_1, \dots, g_{n-1} \in G$ such that the matrix obtained by applying F_n to this collection is lower triangular with integer entries, with positive values on the main diagonal. □

We now prove Theorem 5.1, using the characterization of the commutator subgroup as the elements of compact support. For H_n, the ends correspond to the rays in the obvious way.

Proof Suppose first that G contains the commutator subgroup. Consider the map of coset spaces $\Psi\colon H/G \to \mathbb{Z}^{n-1}/(\boldsymbol{c}\mathbb{Z}^{n-1})$ given by $hG \mapsto F(h)\boldsymbol{c}\mathbb{Z}^{n-1}$. To see that this is well-defined, suppose $h_1G = h_2G$: then $h_1 = h_2g$ for some $g \in G$, and $F(h_1) = F(h_2) + F(g)$. As $F(g) \in \boldsymbol{c}\mathbb{Z}^{n-1}$, we see that $\Psi(h_1G) = \Psi(h_2G)$. We now wish to show that Ψ is injective. Suppose that $\Psi(h_1G) = \Psi(h_2G)$. Then $F(h_1h_2^{-1}) = F(h_1) - F(h_2) = F(g)$ for some $g \in G$. So our question becomes: does $F(h_1h_2^{-1}) \in F(G)$ imply $h_1h_2^{-1} \in G$? The element $h_1h_2^{-1}g^{-1}$ will have no net translation in any end, and hence is compactly supported, and therefore is in G. □

As there are clearly subgroups of finite index in both cases (take the preimage of a finite-index subgroup of $\mathbb{Z}^{n-1}$), there is one loose end remaining: the infinite-index case. Specifically, do there exist subgroups $G < H_n$ (or $G < P\mathcal{H}_n$) whose intersection with the commutator subgroup have infinite index, and such that $\mathrm{fl}(G) = n-1$. First, we see that G is necessarily infinite index in H_n (or $P\mathcal{H}_n$), so all that remains is to check whether this case actually occurs. We thank Noel Brady for the proofs of the following lemmas.

Lemma 5.3 *Suppose $G, K \leq H$ are subgroups. If G is finite index in H, then $G \cap K$ is finite index in K.*

Proof Suppose $[H : G] = m < \infty$. Write $H = Gh_1 \cup \cdots \cup Gh_m$, as a disjoint union. For each index i such that $K \cap Gh_i \neq \varnothing$, there is some $g_i \in G$ and $k_i \in K$ such that $k_i = h_i g_i$. Then we can write

$$Gh_i = Gg_i^{-1}k_i = Gk_i.$$

Thus,

$$K \cap Gh_i = K \cap Gk_i = Kk_i \cap Gk_i = (K \cap G)k_i.$$

Intersecting this with our original disjoint decomposition for H, we have

$$K = (K \cap G)k_1 \cup \cdots \cup (K \cap G)k_s,$$

where $s \leq m$ is the number of indices for which $K \cap Gh_i \neq \varnothing$. □

Lemma 5.4 *Suppose that a group H has an exhaustion by subgroups, ie there exists a nested chain of subgroups $K_1 \leq K_2 \leq \cdots \leq H$ such that H is the union of all the K_i's. Let $G \leq H$. Then $[H : G]$ is finite if and only if the sequence $([K_i : G \cap K_i])$ is eventually constant, in which case it is the limiting value.*

Proof By the proof of Lemma 5.3, we see that the sequence $([K_i : G \cap K_i])$ is bounded by $[H : G] = m$. It is not hard to see that it is nondecreasing, so we need only show that it can't stabilize below m. Choosing a decomposition of H into G-cosets, say $H = Gh_1 \cup \cdots \cup Gh_m$, we see that there must be some index j such that K_j contains all of $\{h_1, \ldots, h_m\}$. Therefore, the sequence of indices stabilizes at m.

For the converse, consider a decomposition $H = Gh_1 \cup Gh_2 \cup \cdots$. Intersecting this with K_i, we see as before that the index $[K_i : G \cap K_i]$ is the number of indices j such that $Gh_j \cap K_i$ is nonempty. As every h_j must be in some K_i, we see that if $[H : G]$ is infinite, the sequence $([K_i : G \cap K_i])$ goes to infinity. □

We return to our question on the existence of infinite-index subgroups of Houghton groups with finiteness length $n-1$, We must now split off the ordinary Houghton group from the surface Houghton group. For H_n, the answer to our existence question is easily yes: there are even copies of H_n itself as infinite-index subgroups of H_n! Consider the stabilizer of a single point in $[n]\times\mathbb{N}$: ignoring the fixed point, and sliding its ray back by one to fill in, we have a natural bijection between $[n]\times\mathbb{N}$, on which H_n acts, and $([n]\times\mathbb{N})\setminus\{(i,j)\}$, on which $\mathrm{Stab}(i,j)$ acts, and we see that these actions are the same. (This fact was known to Houghton; see [15].)

By the same argument, the subgroup fixing pointwise any finite set will be a copy of H_n, and the subgroup fixing (as a set, not pointwise) any finite set will be a finite extension of H_n.

In [10], Yves Cornulier defined a stronger failure of co-Hopfianness, which he called "dis-cohopfian". Call a group G *dis-co-Hopfian* if there is some injective homomorphism $\eta\colon G\to G$ such that the intersection of all iterated images of this homomorphism is trivial, ie

$$\bigcap_{n=1}^{\infty}\eta^n(G)=\{e\}.$$

By taking the finite set being stabilized to be the first point in each $\mathbb{N}$-ray, we obtain the following:

Proposition 5.5 *The Houghton groups H_n are not co-Hopfian, and are in fact dis-co-Hopfian.*

To extend the result on the existence of infinite-index subgroups of finiteness length $n-1$ to surface Houghton groups, we shall work out more carefully the embedding $\mathrm{Br}\,H_n\hookrightarrow\mathcal{H}_n$ suggested by [2]. For the braided Houghton group, we take the asymptotically rigid mapping class group of the following surface, which we shall call $\dot{\Sigma}$ (see [12] for details). Begin with a $2n$-gon as the center piece, and take a punctured square for the attached pieces (see Figure 4). In [13], it was shown that the braided Houghton group $\mathrm{Br}\,H_n$ is type F_{n-1} but not type FP_n. This makes it a good candidate for an infinite-index subgroup of $\mathcal{H}_n$ with $\mathrm{fl}(\mathrm{Br}\,H_n)=\mathrm{fl}(\mathcal{H}_n)$.

To obtain a homomorphism $\mathrm{Br}\,H_n\to\mathcal{H}_n$, we replace each puncture with a boundary component, and then pass to the double. Write Σ' for the surface obtained by replacing punctures with boundary components. As for why this yields an injective homomorphism, consider the following diagram, where the upper row is exact:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & K & \longrightarrow & \mathrm{Map}(\Sigma') & \xrightarrow{a} & \mathrm{Map}(\dot{\Sigma}) & \longrightarrow & 1 \\ & & & & \downarrow{\scriptstyle b} & & & & \\ & & & & \mathrm{Map}(\Sigma) & & & & \end{array}$$

The failure of b to be an injective homomorphism is precisely the subgroup generated by boundary parallel twists. The kernel of a is generated by twists parallel to the compact boundary components. Therefore, we obtain an injective homomorphism $\mathrm{Map}(\dot{\Sigma})\to\mathrm{Map}(\Sigma)$. Restricting to asymptotically rigid subgroups

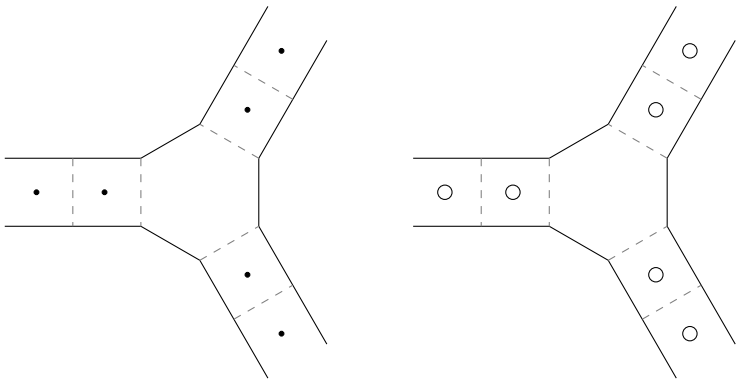

Figure 4: Left: the defining surface for Br H_n. Right: the surface Σ'.

yields an injective homomorphism Br $H_n \to \mathcal{H}_n$. As Br H_n acts trivially on the maximal (ie nonpuncture) ends, this image lies in $P\mathcal{H}_n$.

The image of this homomorphism is a subgroup which fixes (up to isotopy) the multicurve defined by the images of the curve β in Figure 5 in each piece, via the canonical maps ι_B of Section 2.1.

All that remains is to see that this subgroup has infinite index in $P\mathcal{H}_n$. Consider its intersection with the mapping class group of any suited subsurface: here, it must fix the multicurve consisting of the above curves, and is therefore of infinite index. As infinite index in a subgroup implies infinite index in the full group, we have our result.

Between Theorem 5.1, Proposition 5.5, and the above discussion, we have proven the following:

Theorem 5.6 *Let H denote either the Houghton group, or the pure surface Houghton group, and suppose $G < H$ has $\mathrm{fl}(G) = \mathrm{fl}(H)$. Then G is finite index in H if and only if $G \cap H'$ is finite index in H', where H' denotes the commutator subgroup of H. Furthermore, there exist subgroups G with $\mathrm{fl}(G) = \mathrm{fl}(H)$ of both finite and infinite index.*

We finish with further consideration of co-Hopfianness. It is easy to see that Br H_n is not co-Hopfian: there is an inclusion map from the defining surface of Br H_n to itself, sliding the first puncture in some end out, and with all elements moving things only on one side of the skipped puncture (see Figure 6).

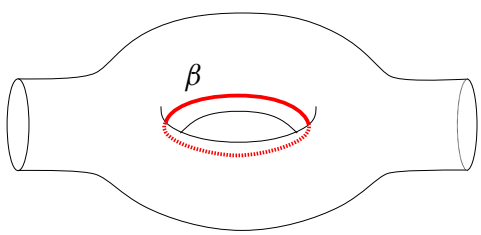

Figure 5: The central curve β in a piece which is fixed by the image of Br H_n in $P\mathcal{H}_n$.

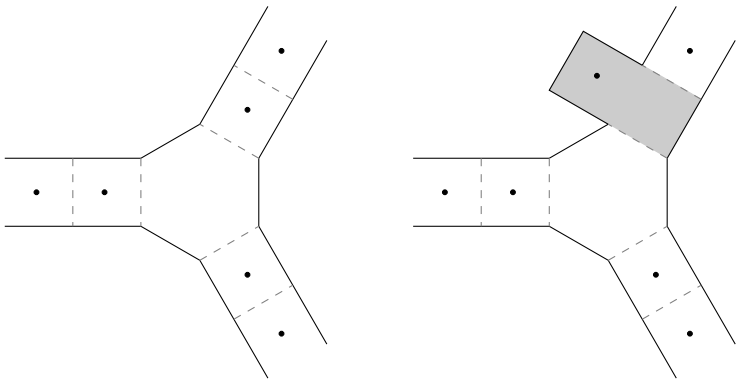

Figure 6: Pushing one puncture "off to the side".

As with the Houghton groups, doing such a move in all ends simultaneously yields a homomorphism whose iterated images act trivially on arbitrarily large compact subsurfaces. This yields:

Theorem 5.7 *The braided Houghton group* Br H_n *is not co-Hopfian, and is in fact dis-co-Hopfian.*

These approaches are not immediately available for the surface Houghton group, as there is no inclusion of surfaces which skips over a single genus. In fact, in light of the results of [4], if the pure surface Houghton group were to fail to be co-Hopfian, then it must fail either by a non-twist-preserving homomorphism, or by a homomorphism which is not the restriction of a homomorphism on the level of pure mapping class groups. This leaves us with the following question:

Question 5.8 *Is the pure surface Houghton group $P\mathcal{H}_n$ co-Hopfian? If not, is it dis-co-Hopfian?*

References

[1] **M Abadie**, CAT(0) *cube complexes and asymptotically rigid mapping class groups*, master's thesis, École Polytechnique Fédérale de Lausanne (2024)

[2] **J Aramayona**, **K-U Bux**, **H Kim**, **C J Leininger**, *Surface Houghton groups*, Math. Ann. 389 (2024) 4301–4318 MR

[3] **J Aramayona**, **G Domat**, **C J Leininger**, *Isomorphisms and commensurability of surface Houghton groups*, J. Group Theory 27 (2024) 1129–1141 MR

[4] **J Aramayona**, **C J Leininger**, **A McLeay**, *Big mapping class groups and the co-Hopfian property*, Michigan Math. J. 74 (2024) 253–281 MR

[5] **J Aramayona**, **P Patel**, **N G Vlamis**, *The first integral cohomology of pure mapping class groups*, Int. Math. Res. Not. 2020 (2020) 8973–8996 MR

[6] **R Bieri**, **W D Neumann**, **R Strebel**, *A geometric invariant of discrete groups*, Invent. Math. 90 (1987) 451–477 MR

[7] **R Bieri**, **B Renz**, *Valuations on free resolutions and higher geometric invariants of groups*, Comment. Math. Helv. 63 (1988) 464–497 MR

[8] **K S Brown**, *Finiteness properties of groups*, J. Pure Appl. Algebra 44 (1987) 45–75 MR

[9] **V Chepoi**, *Graphs of some* CAT(0) *complexes*, Adv. Appl. Math. 24 (2000) 125–179 MR

[10] **Y Cornulier**, *Gradings on Lie algebras, systolic growth, and cohopfian properties of nilpotent groups*, Bull. Soc. Math. France 144 (2016) 693–744 MR

[11] **F Degenhardt**, *Endlichkeitseigenschaften gewisser Gruppen von Zöpfen unendlicher Ordnung*, PhD thesis, Goethe-Universität zu Frankfurt am Main (2000)

[12] **L Funar**, *Braided Houghton groups as mapping class groups*, An. Ştiinţ. Univ. Al. I. Cuza Iaşi. Mat. 53 (2007) 229–240 MR

[13] **A Genevois**, **A Lonjou**, **C Urech**, *Asymptotically rigid mapping class groups*, *II*: *Strand diagrams and nonpositive curvature* (2021) arXiv 2110.06721 To appear in Trans. Amer. Math. Soc.

[14] **A Genevois**, **A Lonjou**, **C Urech**, *Asymptotically rigid mapping class groups, I*: *Finiteness properties of braided Thompson's and Houghton's groups*, Geom. Topol. 26 (2022) 1385–1434 MR

[15] **C H Houghton**, *The first cohomology of a group with permutation module coefficients*, Arch. Math. (Basel) 31 (1978) 254–258 MR

[16] **S R Lee**, *Geometry of Houghton's groups*, PhD thesis, University of Oklahoma (2012) Available at https://www.proquest.com/docview/1044378650

[17] **P Patel**, **N G Vlamis**, *Algebraic and topological properties of big mapping class groups*, Algebr. Geom. Topol. 18 (2018) 4109–4142 MR

[18] **B Renz**, *Geometrische Invarianten und Endlichkeitseigenschaften von Gruppen*, PhD thesis, Goethe-Universität zu Frankfurt am Main (1988) Available at https://esb-dev.github.io/mat/diss.pdf

[19] **M C B Zaremsky**, *On the Σ-invariants of generalized Thompson groups and Houghton groups*, Int. Math. Res. Not. 2017 (2017) 5861–5896 MR

[20] **M C B Zaremsky**, *The BNSR-invariants of the Houghton groups, concluded*, Proc. Edinb. Math. Soc. 63 (2020) 1–11 MR

Department of Mathematics, University of Oklahoma
Norman, OK, United States

Department of Mathematics, University of Oklahoma
Norman, OK, United States

nmtorger@ou.edu, jeremy.west-1@ou.edu

Received: 20 March 2024 Revised: 7 June 2024

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Kathryn Hess
kathryn.hess@epfl.ch
École Polytechnique Fédérale de Lausanne

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Thomas Koberda	University of Virginia thomas.koberda@virginia.edu
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Markus Land	LMU München markus.land@math.lmu.de
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Octav Cornea	Université' de Montreal cornea@dms.umontreal.ca	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
David Futer	Temple University dfuter@temple.edu	Vesna Stojanoska	Univ. of Illinois at Urbana-Champaign vesna@illinois.edu
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Kristen Hendricks	Rutgers University kristen.hendricks@rutgers.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com
Daniel Isaksen	Wayne State University isaksen@math.wayne.edu		

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2025 is US $760/year for the electronic version, and $1110/year (+$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

© 2025 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 25 Issue 8 (pages 4437–5174) 2025

Hierarchies for relatively hyperbolic virtually special groups 4437

EDUARD EINSTEIN

Intersection norms on surfaces and Birkhoff sections for geodesic flows 4499

MARCOS COSSARINI and PIERRE DEHORNOY

Thin knots and the cabling conjecture 4547

ROBERT DEYESO III

Linear linkless embeddings: proof of a conjecture by Sachs 4585

LYNN STANFIELD

Constructing rational homology 3-spheres that bound rational homology 4-balls 4599

LISA LOKTEVA

Homological stability for the ribbon Higman–Thompson groups 4633

RACHEL SKIPPER and XIAOLEI WU

A group-theoretic framework for low-dimensional topology, or: how not to study low-dimensional topology? 4667

SARAH BLACKWELL, ROBION KIRBY, MICHAEL KLUG, VINCENT LONGO and BENJAMIN RUPPIK

Classification of genus-two surfaces in S^3 4719

FILIPPO BARONI

Meromorphic projective structures: signed spaces, grafting and monodromy 4787

SPANDAN GHOSH and SUBHOJOY GUPTA

A new twist on modular links from an old perspective 4827

KHANH LE

Flat fully augmented links are determined by their complements 4839

CHRISTIAN MILLICHAP and ROLLAND TRAPP

BNSR-invariants of surface Houghton groups 4897

NOAH TORGERSON and JEREMY WEST

Topological symmetry groups of the generalized Petersen graphs 4921

ANGELYNN ÁLVAREZ, ERICA FLAPAN, MARK HUNNELL, JOHN HUTCHENS, EMILLE LAWRENCE, PAUL LEWIS, CANDICE PRICE and RUTH VANDERPOOL

Crushing surfaces of positive genus 4949

BENJAMIN A BURTON, THIAGO DE PAIVA, ALEXANDER HE and CONNIE ON YU HUI

Annular links from Thompson's group T 5013

LOUISA LILES

Realizing pairs of multicurves as cylinders on translation surfaces 5031

JULIET AYGUN, JANET BARKDOLL, AARON CALDERON, JENAVIE LORMAN and THEODORE SANDSTROM

Involutive Khovanov homology and equivariant knots 5059

TAKETO SANO

A diagrammatic computation of abelian link invariants 5113

DAVID CIMASONI, LIVIO FERRETTI and JESSICA LIU

Equivariant double-slice genus, stabilization, and equivariant stabilization 5137

MALCOLM GABBARD

Coarse and bi-Lipschitz embeddability of subspaces of the Gromov–Hausdorff space into Hilbert spaces 5153

NICOLÒ ZAVA