THE INTEGRAL COHOMOLOGY OF THE GROUP OF LOOPS

CRAIG JENSEN, JON MCCAMMOND, AND JOHN MEIER

1. INTRODUCTION

Let L_n be a collection of n unknotted, unlinked circles in 3-space, and let $\mathrm{P}\Sigma_n$ be the group of motions of L_n where each circle ends up back at its original position. This group was introduced in the PhD thesis of David Dahm, a student of Ralph Fox, and was later studied by various authors, notably Deborah Goldsmith. Alan Brownstein and Ronnie Lee succeeded in computing $H^2(\mathrm{P}\Sigma_n, \mathbb{Z})$ in [6], and at the end of their paper conjecture a presentation for the algebra $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$. Further evidence for this conjecture came when the cohomological dimension was computed ($\mathrm{cd}(\mathrm{P}\Sigma_n) = n-1$) by Collins in [8], and when the Euler characteristic was computed ($\chi(\mathrm{P}\Sigma_n) = (1-n)^{n-1}$) in [13] (see also [11]). Here we establish the Brownstein-Lee Conjecture. As our argument is a mixture of spectral sequences and combinatorial identities, it seems that Birman was quite prescient in her *Mathematical Review* of the Brownstein-Lee paper: "The combinatorics of the cohomology ring appears to be rich, and the attendant geometric interpretations are very pleasing."

Because $\pi_1(S^3 \setminus L_n) \cong \mathbb{F}_n$, it is not surprising that there is a map

$$\mathrm{P}\Sigma_n \to \textsc{Aut}[\pi_1(S^3 \setminus L_n)] \cong \textsc{Aut}(\mathbb{F}_n) \ .$$

Less immediate is that this map is injective, hence the group $\mathrm{P}\Sigma_n$ can be represented as a group of free group automorphisms. Its image in $\textsc{Out}(\mathbb{F}_n)$ is denoted $\mathrm{OP}\Sigma_n$. (This and other background information is given in §2.) The bulk of our work focuses on the quotient $\mathrm{OP}\Sigma_n$. We compute its integral cohomology groups using the equivariant spectral sequence for the action of $\mathrm{OP}\Sigma_n$ on a contractible complex introduced by McCullough and Miller in [15]. As is often the case, the first page of this spectral sequence is charming yet opaque. Combinatorial arguments are used to show that the E_2 page of the spectral sequence is concentrated in a single column, hence one can read off the cohomology groups from this page. From this we get our Main Theorem and then derive the Brownstein-Lee conjecture.

Main Theorem. *The Poincaré series for $H^*(\mathrm{OP}\Sigma_n, \mathbb{Z})$—a formal power series where the coefficient of z^k is the rank of $H^k(\mathrm{OP}\Sigma_n, \mathbb{Z})$—is $\mathfrak{p}(z) = (1+nz)^{n-2}$.*

Corollary (The Brownstein-Lee Conjecture)**.** *The cohomology of $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$ is generated by one-dimensional classes α^*_{ij} where $i \neq j$, subject to the relations*

1. $\alpha^*_{ij} \wedge \alpha^*_{ij} = 0$
2. $\alpha^*_{ij} \wedge \alpha^*_{ji} = 0$
3. $\alpha^*_{kj} \wedge \alpha^*_{ji} = (\alpha^*_{kj} - \alpha^*_{ij}) \wedge \alpha^*_{ki}$

and the Poincaré series is $\mathfrak{p}(z) = (1+nz)^{n-1}$.

Date: March 8, 2006.

2000 *Mathematics Subject Classification.* 20J06, 57M07.

Remark 1. In the late 1990s, Bogley and Krstić constructed a $K(\mathrm{P}\Sigma_n, 1)$ whose universal cover embeds in Culler and Vogtmann's Outer Space. Using this space they were able to establish the Brownstein-Lee conjecture, but regrettably this work has not appeared.

Remark 2. There are prior results on the (asymptotic) cohomology of $\mathrm{P}\Sigma_n$. In [4] the cohomology of $\mathrm{OP}\Sigma_n$ and $\mathrm{P}\Sigma_n$ with group ring coefficients are determined, and as a corollary, it is shown that $\mathrm{OP}\Sigma_n$ and $\mathrm{P}\Sigma_n$ are duality groups. The ℓ^2-cohomology is computed in [13], where it is shown that the ℓ^2-cohomology of $\mathrm{OP}\Sigma_n$ and $\mathrm{P}\Sigma_n$ is non-trivial and concentrated in top dimension. Alexandra Pettet has recently posted an article explaining the complexity of the kernel of the natural "forgetful" map $\mathrm{P}\Sigma_n \twoheadrightarrow \mathrm{P}\Sigma_{n-1}$ [16].

Acknowledgments. We thank Ethan Berkove, Benson Farb, Allen Hatcher and Alexandra Pettet for their interest and insights into this work. We particularly thank Fred Cohen for sharing some of his work with Jon Pakianathan on an interesting subgroup of $\mathrm{P}\Sigma_n$, and for pointing out a number of connections between this work and results in the literature. (See the end of §6.)

2. $\mathrm{P}\Sigma_n$, $\mathrm{OP}\Sigma_n$ and the complex MM_n

This section moves at a brisk pace. The reader completely unfamiliar with the groups $\mathrm{P}\Sigma_n$ and the McCullough-Miller complexes MM_n should perhaps read the first six sections of [13] where the material summarized in this section is developed in greater detail.

On the intuitive level, $\mathrm{P}\Sigma_n$ is the group of motions of n unknotted, unlinked circles in the 3-sphere. In order to be more precise and efficient, we present $\mathrm{P}\Sigma_n$ as a subgroup of the automorphism group of a free group $\mathbb{F}_n$ (for background on this isomorphism, see [9]). The group of *pure symmetric automorphisms* of $\mathbb{F}_n$ consists of all automorphisms that, for a fixed basis $\{x_1, \ldots, x_n\}$, send each x_i to a conjugate of itself. This group is generated by the automorphisms α_{ij} induced by

$$\alpha_{ij} = \begin{cases} x_i \to x_j x_i x_j^{-1} & \\ x_k \to x_k & k \neq i \end{cases} .$$

McCool proved that the relations

$$\left\{ \begin{array}{ll} [\alpha_{ij}, \alpha_{kl}] & i,j,k \text{ and } l \text{ all distinct} \\ [\alpha_{ij}, \alpha_{kj}] & i,j \text{ and } k \text{ distinct} \\ [\alpha_{ij}, \alpha_{ik}\alpha_{jk}] & i,j \text{ and } k \text{ distinct} \end{array} \right\}$$

are sufficient to present $\mathrm{P}\Sigma_n$ [14]. Given $j \in [n]$ and $I \subset [n] \setminus \{j\}$ we let α_{Ij} denote the product of generators

$$\alpha_{Ij} = \prod_{i \in I} \alpha_{ij} .$$

Using McCool's relations one sees that this product is independent of the order in which one lists the α_{ij}. Note that when $I = [n] \setminus \{j\}$, the element α_{Ij} is simply conjugation by x_j.

For the remainder of this paper we view $\mathrm{P}\Sigma_n$ as a subgroup of $\mathrm{Aut}(\mathbb{F}_n)$. Since the inner automorphisms, $\mathrm{Inn}(\mathbb{F}_n)$, form a subgroup of $\mathrm{P}\Sigma_n$, we may form the quotient $\mathrm{P}\Sigma_n/\mathrm{Inn}(\mathbb{F}_n)$, which we denote $\mathrm{OP}\Sigma_n$. Interestingly, $\mathrm{P}\Sigma_n$ is a subgroup of the famous $\mathrm{IA}_n = \mathrm{Ker}[\mathrm{Aut}(\mathbb{F}_n) \twoheadrightarrow GL_n(\mathbb{Z})]$. In fact, the set $\{\alpha_{ij}\}$ is a subset

of the standard generating set of IA_n discovered by Magnus. The image of $\mathrm{P}\Sigma_n$ contains the image of the pure braid group, P_n, under Artin's embedding, hence there are proper inclusions

$$P_n < \mathrm{P}\Sigma_n < \mathrm{IA}_n < \mathrm{Aut}(\mathbb{F}_n).$$

Viewed as a group of free group automorphisms, there is an action of $\mathrm{P}\Sigma_n$ on a contractible complex constructed by McCullough and Miller [15]. Our perspective on this complex is that of [13] where it is described in terms of marked, $[n]$-labelled hypertrees. We quickly recall the definition of the $[n]$-labelled hypertree poset.

Definition 2.1. A *hypergraph* Γ consists of a set of vertices V and a set of hyperedges E, each element of E containing at least two vertices. We refer to edges in a hypergraph which contain more than two vertices as *fat* edges. A *walk* in a hypergraph Γ is a sequence $v_0, e_1, v_1, \ldots, v_{n-1}, e_n, v_n$ where for all i, $v_i \in V$, $e_i \in E$ and for each e_i, $\{v_{i-1}, v_i\} \subset e_i$. A hypergraph is *connected* if every pair of vertices is joined by a walk. A *simple cycle* is a walk that contains at least two edges, all of the e_i are distinct and all of the v_i are distinct except $v_0 = v_n$. A hypergraph with no simple cycles is a *hyperforest* and a connected hyperforest is a *hypertree*. Note that the no simple cycle condition implies that distinct edges in Γ have at most one vertex in common

An $[n]$-*labelled* hypertree is a hypertree whose vertices have been labelled (bijectively) by $[n] = \{1, \ldots, n\}$. Examples of $[n]$-labelled hypertrees can be found in most figures in this paper. The *rank* of a hypertree τ is $\#E(\tau) - 1$.

For any fixed value of n, one can define a partial order on the set of all $[n]$-labelled hypertrees: $\tau \leq \tau'$ if every edge of τ' is contained in an edge of τ. The poset consisting of $[n]$-labelled hypertrees with this partial ordering is denoted HT_n and is called the *hypertree poset*. In Figure 1 we show a maximal chain in HT_6. Notice that maximal elements in HT_n correspond to ordinary, that is simplicial, trees on $[n]$.

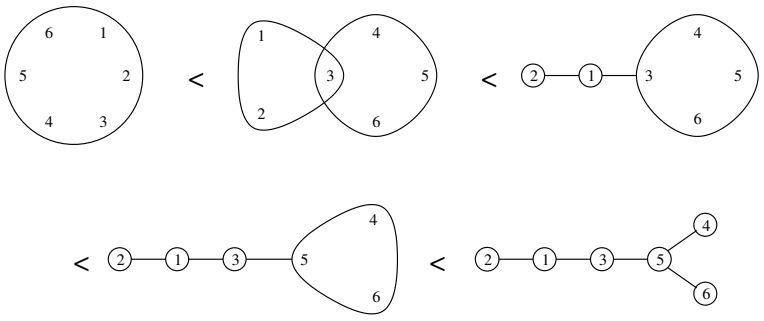

FIGURE 1. A maximal chain in HT_6.

The combinatorics of HT_n are interesting (see [4], [11] and [13]), but by and large we will need few previously established combinatorial facts about this poset. One fact that we will use is:

Lemma 2.2. *Given a collection of hypertrees $\{\tau_1, \ldots, \tau_k\}$ in HT_n there is a unique, maximal hypertree τ such that $\tau \leq \tau_i$ for each i. In other words, HT_n is a meet semi-lattice.*

The McCullough-Miller complex MM_n was introduced in [15] for the study of certain automorphism groups of free products. The main results we need in this paper are summarized in the following Theorem.

Theorem 2.3 ([15])**.** *The space MM_n is a contractible simplicial complex admitting a simplicial $\mathrm{OP}\Sigma_n$ action such that:*

1. *The fundamental domain is a strong fundamental domain and the quotient $\mathrm{OP}\Sigma_n \backslash MM_n$ is isomorphic to the geometric realization of the poset HT_n.*
2. *If σ is a simplex of $|\mathrm{HT}_n|$ corresponding to a chain $\tau_0 < \tau_1 < \cdots < \tau_k$, then*
$$Stab(\sigma) = Stab(\tau_0) \cong \mathbb{Z}^{rk(\tau_0)}$$
3. *Let $j \in [n]$ and let I be a subset of $[n]$ corresponding to all the labels of a connected component of τ minus the vertex labelled j. Then the outer-automorphism given by the automorphism α_{Ij} is contained in $Stab(\tau)$ and the collection of all such outer-automorphisms is a generating set for $Stab(\tau)$.*

The generating set given above is not a minimal generating set. Consider for example the maximal hypertree shown in Figure 1. If we let $\bar{\alpha}$ denote the image of $\alpha \in \mathrm{P}\Sigma_n$ in $\mathrm{OP}\Sigma_n$, then we have $\bar{\alpha}_{45}, \bar{\alpha}_{65}$ and $\bar{\alpha}_{\{1,2,3\},5}$ are all in $\mathrm{Stab}(\tau)$. But
$$\bar{\alpha}_{45} \cdot \bar{\alpha}_{65} = \left[\bar{\alpha}_{\{1,2,3\},5}\right]^{-1} \text{ in } \mathrm{OP}\Sigma_6$$
as $\alpha_{45} \cdot \alpha_{65} \cdot \alpha_{\{1,2,3\},5}$ is an inner automorphism.

Using the equivariant spectral sequence, applied to the action $\mathrm{OP}\Sigma_n \curvearrowright \mathrm{MM}_n$, we compute the cohomology groups $H^i(\mathrm{OP}\Sigma_n, \mathbb{Z})$. Recall that the equivariant spectral sequence for a group G acting simplicially on a contractible complex is given by
$$E_1^{pq} = \prod_{\sigma \in \mathcal{E}_p} H^q(G_\sigma, M) \Rightarrow H^{p+q}(G, M)$$
where $\mathcal{E}_p$ denotes a set of representatives of the G-orbits of p-cells. The differentials on the E_1 page are the standard ones of the equivariant cohomology spectral sequence, namely a combination of restriction maps to a subgroup and coboundary maps. (See §VII.7 of [5].) Since the simplex stabilizers for the action of $\mathrm{OP}\Sigma_n \curvearrowright \mathrm{MM}_n$ are free abelian, and our set of representatives of $\mathrm{OP}\Sigma_n$-orbits can be taken to be p-simplices in the geometric realization $|\mathrm{HT}_n|$, we can exhibit the first page of the equivariant spectral sequence in a fairly concrete manner.

For example, in Figure 2 we show the non-zero portion of the E_1 page for the action of $\mathrm{OP}\Sigma_4$, where we have suppressed the $\mathbb{Z}$-coefficients. The geometric realization of HT_4 was worked out as an example in [15]. We redraw their figure in Figure 3. The actual geometric realization has dimension 2, but the vertex corresponding to the hypertree with exactly one edge forms a cone point in $|\mathrm{HT}_4|$ and so it is not shown in the figure. The left edge of the E_1 page is explained by the following observations: The $\mathrm{OP}\Sigma_4$ orbits of vertices under the action $\mathrm{OP}\Sigma_4 \curvearrowright \mathrm{MM}_4$ correspond to the elements of HT_4, and by direct observation one sees:

- there is one hypertree in HT_4 whose stabilizer is trivial. It is the hypertree with a single hyperedge, not shown in Figure 3.

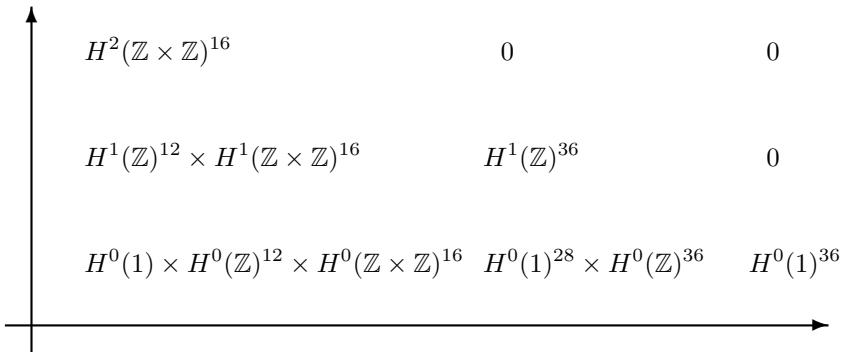

FIGURE 2. The E_1 page of the equivariant spectral sequence for the action $\mathrm{OP}\Sigma_4 \curvearrowright \mathrm{MM}_4$.

- There are twelve hypertrees whose stabilizers are $\cong \mathbb{Z}$. These all have the same combinatorial type, shown in the top right of Figure 3.
- The remaining sixteen hypertrees have stabilizers $\cong \mathbb{Z}\times\mathbb{Z}$.

The interested reader may use Figure 3 to double-check our entries in Figure 2.

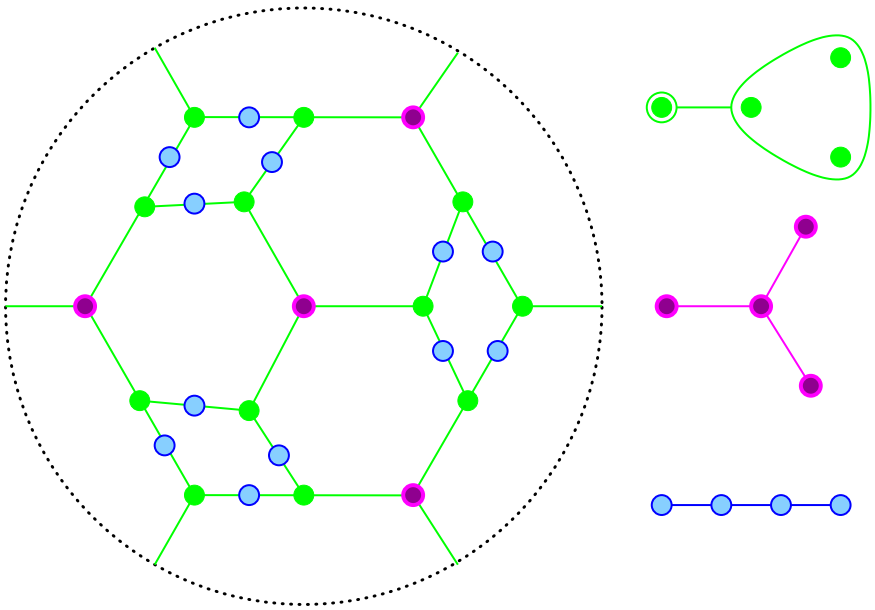

FIGURE 3. The geometric realization of HT_4, excluding the vertex corresponding to a hypertree with a single edge. The graph is to be thought of as embedded in a projective plane, so edges leaving the dotted circle re-enter at the opposite point. The color coding is by combinatorial type.

While we do give a concrete description of the E_1 page of this spectral sequence, the reader should not be lulled into thinking that in general this page is always

directly accessible. For example, Cayley's formula states that the number of (simplicial) trees on $[n]$ is n^{n-2}. This implies that the $(0, n-2)$-entry on the E_1 page for the action $\mathrm{OP}\Sigma_n \curvearrowright \mathrm{MM}_n$ is $H^{n-2}(\mathbb{Z}^{n-2})^{n^{n-2}}$ and this is most certainly not the largest entry on the page.

3. Choosing a Basis

We are able to understand the contents of the E_1 page of the equivariant spectral sequence by finding a concrete, minimal set of generators for each stabilizer of each simplex in $|\mathrm{HT}_n|$. Recall that the stabilizer of a simplex is the stabilizer of its minimal vertex, hence it suffices to pick generators for the stabilizers of individual hypertrees in the poset. Further, the stabilizer of a hypertree τ under the action of $\mathrm{OP}\Sigma_n$ is free abelian of rank $\mathrm{rk}(\tau)$. As an example, the stabilizer under the action of $\mathrm{OP}\Sigma_6$ of the hypertree with two edges in Figure 1 is infinite cyclic. It is generated by the image of $\alpha_{\{1,2\},3}$ in $\mathrm{OP}\Sigma_6$, or equivalently by the image of $\alpha_{\{4,5,6\},3}$.

For a fixed hypertree τ the chosen set of generators is essentially described by the following conditions: You pick generators as listed in Theorem 2.3 except,

1. You never conjugate x_1, and
2. You never conjugate x_2 by x_1.

We now make this precise.

Definition 3.1. For each $j \in [n]$ let $\hat{j}$ be the index that is to be avoided when conjugating by x_j. That is, $\hat{j} = 1$ for $j \neq 1$ and $\hat{1} = 2$. Define a *one-two automorphism* to be any automorphism of the form $\alpha_{Ij} = \prod_{i \in I} \alpha_{ij}$, where $I \subset \{[n] \setminus \{j, \hat{j}\}\}$. The image of α_{Ij} in $\mathrm{OP}\Sigma_n$ is denoted $\bar{\alpha}_{Ij}$ and is called a *one-two outer-automorphism*.

For a given hypertree τ the collection of all one-two outer-automorphisms in $\mathrm{Stab}(\tau)$ is the *one-two basis* for $\textsc{Stab}(\tau)$. We denote this by $\mathcal{B}(\tau)$. If σ is a simplex in $|\mathrm{HT}_n|$ corresponding to a chain $\tau_0 < \tau_1 < \cdots < \tau_k$ then the stabilizer of σ is the stabilizer of τ_0 and so we define $\mathcal{B}(\sigma) = \mathcal{B}(\tau_0)$. A simplex σ in $|\mathrm{HT}_n|$ is in the *support* of a one-two outer-automorphism $\bar{\alpha}_{Ij}$ if $\bar{\alpha}_{Ij} \in \mathcal{B}(\sigma)$.

If one roots a hypertree τ at the vertex labelled 1, then the one-two basis for τ can be viewed as being (partly) induced by "gravity". The prescription to avoid conjugating x_1 by an x_j corresponds to having x_j conjugate the elements corresponding to vertices below the vertex labelled j. The exception is if the vertex labelled 1 has valence greater than one, in which case it conjugates the elements below it, but not those in the branch containing the vertex labelled by 2. Thus, the one-two basis for the stabilizer of the hypertree shown in Figure 4 is

$$\mathcal{B}(\tau) = \{\bar{\alpha}_{\{4,5,8,9,10,11\},1}, \bar{\alpha}_{\{6,7\},2}, \bar{\alpha}_{\{8,9\},4}, \bar{\alpha}_{\{10,11\},5}\}$$

The size of a one-two basis depends only on the rank of the underlying hypertree, but the actual elements in the one-two basis very much depend on the labelling. Two rank-two, labelled hypertrees are shown in Figure 5. The underlying hypertrees are the same, but the labels of the vertices are different. Both have stabilizers isomorphic to $\mathbb{Z} \times \mathbb{Z}$. The one-two generating set for the stabilizer of the labelled hypertree on the left is $\{\bar{\alpha}_{\{5,6\},1}, \bar{\alpha}_{\{2,3\},4}\}$, while the generating set for the stabilizer of the hypertree on the right is $\{\bar{\alpha}_{\{2,4,5\},3}, \bar{\alpha}_{\{2,4\},5}\}$.

Finally, one last example: The stabilizer of the highest rank hypertree in Figure 1 is $\cong \mathbb{Z}^4$ and the one-two generating set for $\textsc{Stab}(\tau)$ is

$$\mathcal{B}(\tau) = \{\bar{\alpha}_{\{3,4,5,6\},1}, \bar{\alpha}_{\{4,5,6\},3}, \bar{\alpha}_{4,5}, \bar{\alpha}_{6,5}\} \ .$$

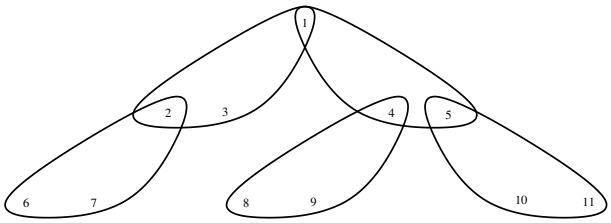

FIGURE 4. Rooting a hypertree at the vertex labelled 1.

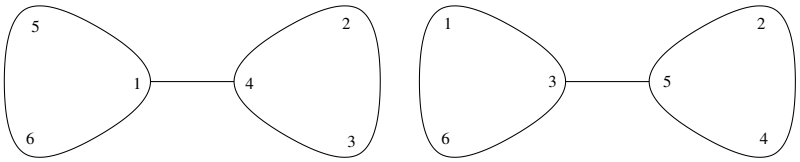

FIGURE 5. Two rank two labelled hypertrees in HT_6.

Proposition 5.1 of [15] gives a concrete description of the stabilizers of vertices, from which one can derive:

Lemma 3.2. *The set $\mathcal{B}(\tau)$ is a minimal rank generating set for Stab(τ).*

(Hence our use of the term "basis".)

A one-two outer-automorphism $\bar{\alpha}_{Ij}$ may be in the stabilizer of a simplex without being part of the one-two basis for the simplex. Consider for example the hypertrees shown in Figure 6. The stabilizer of the rank one hypertree on the left is generated by $\mathcal{B}(\tau) = \{\bar{\alpha}_{\{5,6,7\},4}\}$. However, if τ' is the hypertree on the right, then $\mathcal{B}(\tau') = \{\bar{\alpha}_{54}, \bar{\alpha}_{\{6,7\},4}, \bar{\alpha}_{76}\}$. Thus $\bar{\alpha}_{\{5,6,7\},4} \in \mathrm{Stab}(\tau')$ but $\bar{\alpha}_{\{5,6,7\},4} \notin \mathcal{B}(\tau')$.

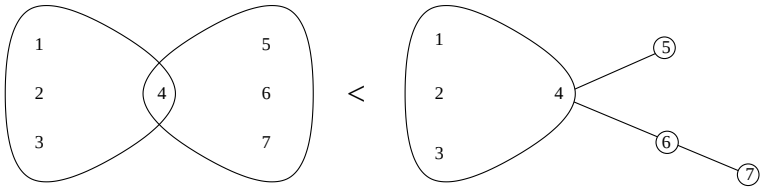

FIGURE 6. The outer-automorphism $\bar{\alpha}_{\{5,6,7\},4}$ is contained in the stabilizer of both hypertrees. It is in the one-two basis of the hypertree on the left, but it is not in the basis of the hypertree on the right. (It is in the basis of the 1-simplex corresponding to the chain $\tau < \tau'$.)

The distinction between being in the stabilizer and being in the one-two basis can be characterized combinatorially by removing vertices. Given any hypertree τ and a vertex j, we can construct a new hypergraph on $[n] \setminus \{j\}$ by simply removing j from each hyperedge set and then removing any singleton sets that result. Let $\bar{\alpha}_{Ij}$ be a one-two outer automorphism, let τ be a hypertree and let τ' be the hyperforest obtained by removing j from the vertex set. The stabilizer of τ contains $\bar{\alpha}_{Ij}$ if and

only if I is a union of vertices of connected components of τ' and the one-two basis of τ contains $\bar{\alpha}_{Ij}$ if and only if I is the vertex set of single connected component of τ'.

Here's the lovely fact. Every simplex in $|\mathrm{HT}_n|$ that supports a one-two outer-automorphism $\bar{\alpha}_{Ij}$ is compatible with the hypertree consisting of two edges—$I \cup \{j\}$ and $[n] \setminus I$—joined along the vertex labelled j. Moreover, the one-two basis for this two edge hypertree is $\{\bar{\alpha}_{Ij}\}$. (See Figure 7.) These facts extend to all subsets of one-two bases (Lemma 3.6).

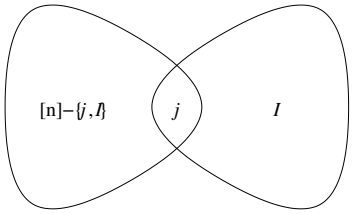

FIGURE 7. The cone point for a one-two basis element.

Definition 3.3. A collection $\mathcal{A}$ of one-two outer-automorphisms is called *compatible* if it is a subset of a one-two basis for some hypertree τ ($\mathcal{A} \subset \mathcal{B}(\tau)$). (The notion of a compatible collection of one-two outer-automorphisms is essentially the same as McCullough and Miller's notion of pairwise disjoint based partitions of $[n]$ in [15].)

Definition 3.4. Let $\mathcal{A}$ be a compatible collection of one-two outer-automorphisms. The collection of hypertrees that support $\mathcal{A}$ is the $\mathcal{A}$-*core*. Call an element of the hypertree poset that contains the set $\mathcal{A}$ in its stabilizer—but does not support it in terms of being part of the one-two basis for the hypertree—an $\mathcal{A}$-*peripheral* hypertree. Further, call the subcomplex induced by the $\mathcal{A}$-core hypertrees the $\mathcal{A}$-*core complex*, and the subcomplex induced by $\mathcal{A}$-peripheral hypertrees the $\mathcal{A}$-*peripheral complex*.

The fact that the hypertree poset has meets (Lemma 2.2) implies the following:

Lemma 3.5. *If $\mathcal{A}$ is a compatible collection of one-two outer-automorphisms then there is a hypertree $\tau(\mathcal{A})$, called the* cone point of $\mathcal{A}$ *such that if $\mathcal{A} \subset \mathrm{STAB}(\tau')$ then $\tau(\mathcal{A}) \leq \tau'$.*

In fact, even more is true. An easy induction on the number of edges in τ (or the size of the compatible collection $\mathcal{A}$, respectively) can be used to generalize the lovely facts listed above.

Lemma 3.6. *For all hypertrees τ, the cone point of the one-two basis of τ is τ itself ($\tau(\mathcal{B}(\tau)) = \tau$) and for all compatible collections $\mathcal{A}$, the one-two basis of the cone point of $\mathcal{A}$ is $\mathcal{A}$ itself ($\mathcal{B}(\tau(\mathcal{A})) = \mathcal{A}$). As a consequence, the cone point hypertree $\tau(\mathcal{A})$ lies in the $\mathcal{A}$-core.*

The combinatorics of the spectral sequence break into two cases.

Definition 3.7. A compatible collection of one-two outer-automorphisms $\mathcal{A}$ is *essential* if $\tau(\mathcal{A})$ satisfies

1. There is at most one fat edge (an edge of order > 2);
2. If there is a fat edge then the vertex labelled 1 is in the fat edge;
3. If there is a fat edge then the unique reduced path from the vertex labelled 1 to the vertex labelled 2 contains the fat edge.

Notice that $\mathcal{A}$ is essential when $\tau(\mathcal{A})$ is an ordinary tree. Compatible collections of one-two outer-automorphisms that are not essential are *inessential*.

For example, $\mathcal{A} = \{\bar{\alpha}_{54}, \bar{\alpha}_{\{6,7\},4}, \bar{\alpha}_{76}\}$ is an essential basis as it is the one-two basis for the stabilizer of the rank three hypertree shown in Fig. 6.

To build some intuition for what can occur with inessential one-two bases, consider the one-two outer-automorphism $\bar{\alpha}_{\{3,4,5\},2} \in \mathrm{OP}\Sigma_5$. The cone point $\tau(\{\bar{\alpha}_{\{3,4,5\},2}\})$ is shown in Figure 8. While there is a single fat edge, the fat edge

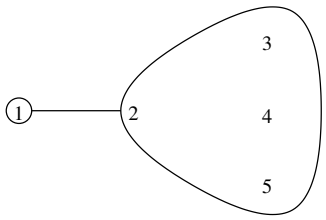

FIGURE 8. The hypertree $\tau(\{\bar{\alpha}_{\{3,4,5\},2}\})$.

does not contain the vertex labelled by 1, nor is it in the minimal path from the vertex labelled 1 to the vertex labelled 2. The associated $\{\bar{\alpha}_{\{3,4,5\},2}\}$-peripheral subcomplex of $|\mathrm{HT}_5|$ is shown in Figure 9. The outer-automorphism $\bar{\alpha}_{\{3,4,5\},2}$ is in the stabilizer of every hypertree in this subcomplex, but it is not in the one-two basis for any of these hypertrees.

Lemma 3.8. *Let $\mathcal{A}$ be a compatible collection of one-two automorphisms and let $\tau(\mathcal{A})$ be the cone point of $\mathcal{A}$. Then $\mathcal{A}$ is essential if and only if the $\mathcal{A}$-peripheral complex is empty.*

Proof. Let $\mathcal{A}$ be essential and let $\alpha_{Ij} \in \mathcal{A}$. The sub-hypertree induced by $I \cup \{j\}$ is an ordinary tree, with the vertex corresponding to j being a leaf. Since every hypertree τ with $\tau(\mathcal{A}) < \tau$ is formed by dividing the single fat edge of $\tau(\mathcal{A})$ into a higher rank hypertree, the sub-hypertree of τ induced by $I \cup \{j\}$ is the same ordinary tree. Hence $\alpha_{Ij} \in \mathcal{B}(\tau)$, so $\mathcal{A} \subset \mathcal{B}(\tau)$. Thus by definition the $\mathcal{A}$-peripheral complex is empty.

If $\mathcal{A}$ is inessential, there must be a one-two automorphism α_{Ij} where the sub-hypertree induced by $I \cup \{j\}$ contains a fat edge, e. Let k label the vertex of e that is closest to j. Then splitting e into two hyperedges that are joined along k creates a hypertree τ that is above $\tau(\mathcal{A})$ in HT_n but $\alpha_{Ij} \notin \mathcal{B}(\tau)$. Thus τ is peripheral, and the $\mathcal{A}$-peripheral complex is not empty. □

4. Computing the E_2 page

In this section we establish:

Proposition 4.1. *The non-zero entries on the E_2 page for the spectral sequence corresponding to the action $\mathrm{OP}\Sigma_n \curvearrowright MM_n$ are concentrated in the $(0, q)$-column.*

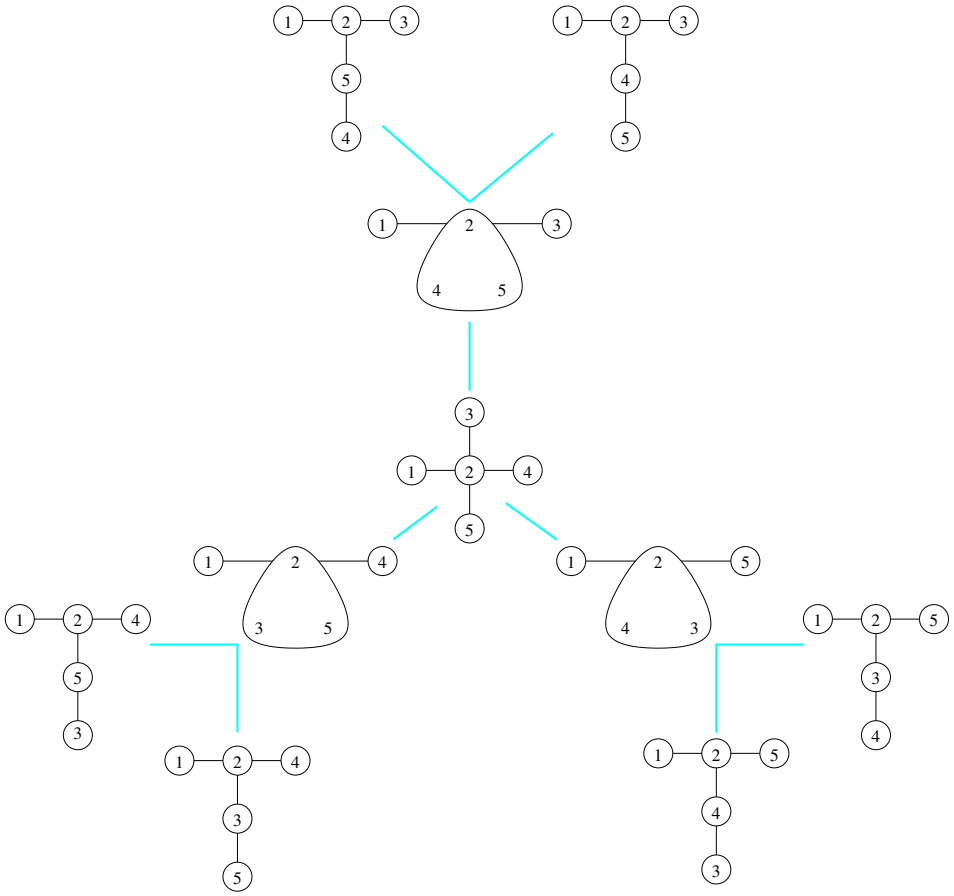

FIGURE 9. The peripheral subcomplex of $|\mathrm{HT}_5|$ associated to the one-two outer-automorphism $\bar{\alpha}_{\{3,4,5\},2}$.

To start the process of proving this, we note that the one-two bases for the stabilizers of hypertrees allow us to give a concrete description of the entries on the first page of the spectral sequence. The (p, q) entry on the E_1 page is the product

$$\prod_{\substack{\sigma \in |\mathrm{HT}_n|, \\ \dim(\sigma) = p}} H^q(\textsc{Stab}(\sigma), \mathbb{Z}) \ .$$

Since the stabilizer of σ is the stabilizer of its minimal element, we may consider what happens in the case of a hypertree τ. Each stabilizer is free abelian, so the first homology $H_1(\mathrm{Stab}(\tau), \mathbb{Z})$ is free abelian with generating set $\{[\bar{\alpha}_{Ij}] \mid \bar{\alpha}_{Ij} \in \mathcal{B}(\tau)\}$. The cohomology group $H^1(\mathrm{Stab}(\tau), \mathbb{Z})$ is then generated by the dual basis $\{\bar{\alpha}^*_{Ij}\}$ where

$$\bar{\alpha}^*_{Ij}([\bar{\alpha}_{Kl}]) = \begin{cases} 1 & I = K \text{ and } j = l \\ 0 & \text{otherwise} \end{cases} \ .$$

Because the $H^*(\mathbb{Z}^n)$ is an exterior algebra generated by one-dimensional classes, the set $\{\bar{\alpha}^*_{Ij} \mid \bar{\alpha}_{Ij} \in \mathcal{B}(\tau)\}$ is a generating set for the cohomology of the stabilizer.

It follows that the set of all products of q distinct $\bar{\alpha}_{Ij}$s is a generating set for $H^q(\textsc{Stab}(\tau), \mathbb{Z})$. Thus we may identify the generating set for $H^q(\textsc{Stab}(\sigma), \mathbb{Z})$ with the collection of all subsets of one-two outer-automorphisms $\mathcal{A} \subset \mathcal{B}(\tau)$ with $|\mathcal{A}| = q$.

As the differentials are a combination of restrictions to direct summands and coboundary maps, the E_1 page of the equivariant spectral sequence can be expressed as a union of sub-cochain complexes. Let $\mathcal{A}$ be a compatible collection of one-two automorphisms, with $|\mathcal{A}| = q$. Then in row q one sees the terms where $\mathcal{A}$ describes a generator, corresponding to simplices σ with $\mathcal{B}(\sigma) \supset \mathcal{A}$.

Lemma 4.2. *If $\mathcal{A}$ is essential, then the collection of all entries in row q of the first page that are given by having $\mathcal{A} \subset \mathcal{B}(\tau)$ forms a sub-cochain complex whose cohomology consists of a single $\mathbb{Z}$ in dimension zero.*

Proof. Since $\mathcal{A}$ is essential, its support is a subcomplex of the hypertree poset. That is, if $\sigma \in |\mathrm{HT}_n|$ is in the support of $\mathcal{A}$, then every face of σ is in the support of $\mathcal{A}$. Further, the vertex associated to $\tau(\mathcal{A})$ is a cone point for the support of $\mathcal{A}$, hence the support of $\mathcal{A}$ is contractible. The fact that the support is actually a subcomplex shows that the sub-cochain complex in the E_1 page that corresponds to $\mathcal{A}$ is just the cochain complex for the support of $\mathcal{A}$. Because the support is contractible, the cohomology of this cochain complex is trivial except for a single $\mathbb{Z}$ in dimension zero. □

We now turn to the case where $\mathcal{A}$ is inessential. Call a fat edge in $\tau(\mathcal{A})$ *worrisome* if it does not contain the vertex labelled 1 or if it does contain the vertex labelled 1, but the reduced path joining 1 to 2 does not contain this edge. Let $\{e_1, \ldots, e_k\}$ be the worrisome fat edges of $\tau(\mathcal{A})$. Let c_i be the label of the vertex in e_i that is closest to the vertex labelled 1 and let L_i be the remaining labels of vertices in e_i. Thus, viewing $\tau(\mathcal{A})$ as a hypertree rooted at 1, the vertices labelled by numbers in L_i are one level down from the vertex c_i.

If $\mathcal{P}_i$ is a partition of L_i define $\tau(\mathcal{A})_{\mathcal{P}_i}$ to be the hypertree where the hyperedge e_i has been split into hyperedges defined by this partition. That is, if $\mathcal{P}_i = \{\ell_1, \ldots, \ell_m\}$ then the edge e_i will be replaced with the edges $\{\ell_j \cup \{c_i\} \mid 1 \leq j \leq m\}$. Moreover, if $\mathfrak{p} \in \Pi_1 \times \cdots \times \Pi_k$ is a collection of partitions of the L_i, define $\tau(\mathcal{A})_{\mathfrak{p}}$ to be the hypertree formed by splitting each worrisome edge e_i by the associated partition of L_i.

Example 4.3. Consider the hypertree τ shown in Figure 4. The worrisome edges are all of the edges of τ, excepting $e = \{1, 2, 3\}$. Set $L_1 = \{4, 5\}, L_2 = \{6, 7\}, L_3 = \{8, 9\}$ and $L_4 = \{10, 11\}$ in order to fix the indexing. Let

$$\mathfrak{p} = (\{\{4\}, \{5\}\}, \{\{6, 7\}\}, \{\{8, 9\}\}, \{\{10\}, \{11\}\})$$

be the result of taking non-trivial partitions of L_1 and L_4. Then the tree $\tau_{\mathfrak{p}}$ formed by splitting τ according to this product of partitions is shown in Figure 10.

Lemma 4.4. *Let $\mathcal{A}$ be an inessential collection of one-two automorphisms. Then the $\mathcal{A}$-peripheral complex is contractible.*

Proof. Let $\mathfrak{p}$ be a non-trivial element of $\Pi_1 \times \cdots \times \Pi_k$. That is, assume at least one of the partitions is non-trivial. Let $\mathrm{HT}_{\geq \mathfrak{p}}$ be the subposet of HT_n consisting of elements greater than or equal to $\tau(\mathcal{A})_{\mathfrak{p}}$. The associated order complexes $|\mathrm{HT}_{\geq \mathfrak{p}}|$ cover the peripheral complex.

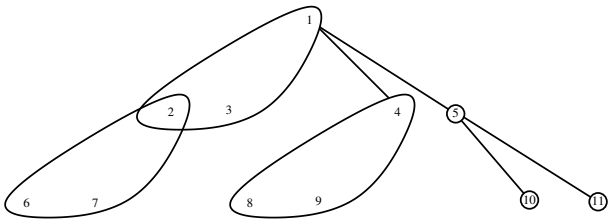

FIGURE 10. The hypertree resulting from splitting some of the worrisome fat edges in the hypertree shown in Figure 4.

Each $|\mathrm{HT}_{\geq \mathfrak{p}}|$ is contractible as the vertex associated to $\tau(\mathcal{A})_{\mathfrak{p}}$ forms a cone point. Further, if $\mathfrak{p}_1, \ldots, \mathfrak{p}_s$ is any collection of non-trivial elements in $\Pi_1 \times \cdots \times \Pi_k$ then their meet $\mathfrak{p} = \mathfrak{p}_1 \wedge \cdots \wedge \mathfrak{p}_s$ is a non-trivial element of $\Pi_1 \times \cdots \times \Pi_k$ and

$$|\mathrm{HT}_{\geq \mathfrak{p}}| = |\mathrm{HT}_{\geq \mathfrak{p}_1}| \cap \cdots \cap |\mathrm{HT}_{\geq \mathfrak{p}_s}| \ .$$

Thus we have covered the peripheral complex by contractible subcomplexes whose intersections are also contractible. By the Quillen Fiber Lemma, the peripheral complex is homotopy equivalent to the nerve of this covering.

Let $\hat{\mathfrak{p}}$ be the product of partitions given by totally partitioning each L_i. Then $|\mathrm{HT}_{\geq \hat{\mathfrak{p}}}|$ is a single vertex that is a cone point in the nerve of the covering. Thus the nerve of the covering is contractible, hence so is the $\mathcal{A}$-peripheral complex. □

The argument above may become less opaque if one consults the example of a peripheral complex given in Figure 9. The original $\tau(\mathcal{A})$ is shown in Figure 8, and the single worrisome edge is $e = \{2, 3, 4, 5\}$. The reader will find that this peripheral complex has been covered by four contractible subcomplexes corresponding to the four non-trivial partitions of $\{3, 4, 5\}$.

Lemma 4.5. *If $\mathcal{A}$ is inessential, then the collection of all entries in row q of the first page that are given by having $\mathcal{A} \subset \mathcal{B}(\sigma)$ forms a sub-cochain complex whose cohomology is trivial in all dimensions.*

Proof. Let $\mathcal{A} \subset \mathcal{B}(\sigma)$. Then σ corresponds to a chain of hypertrees $\tau_0 < \cdots < \tau_p$ and $\mathcal{A} \subset \mathcal{B}(\tau_0)$. Since $\tau(\mathcal{A}) \leq \tau_0$ we can pair p-simplices where $\tau_0 \neq \tau(A)$ with $(p+1)$-simplices formed by adding $\tau(\mathcal{A})$ to the chain:

$$\underbrace{\tau_0 < \cdots < \tau_p}_{\sim \sigma} \leftrightarrow \underbrace{\tau(\mathcal{A}) < \tau_0 < \cdots < \tau_p}_{\sim \sigma'} \ .$$

We can remove the terms corresponding to paired simplices in the original cochain complex to form a cochain complex with equivalent cohomology; in other words we may restrict ourselves to cochains with support on chains with initial element $\tau(\mathcal{A})$. Thus the terms in this new cochain complex correspond to chains of hypertrees of the form $\tau(\mathcal{A}) < \tau_1 < \cdots < \tau_k$ where each τ_i (for $i \geq 1$) is in the $\mathcal{A}$-peripheral subcomplex of $|\mathrm{HT}_n|$. In dimension zero there is a single $\mathbb{Z}$ corresponding to $\tau(\mathcal{A})$. In general, in dimension k there is a $\mathbb{Z}^{p(k)}$ if $p(k)$ is the number of $(k-1)$-simplices in the $\mathcal{A}$-peripheral subcomplex. Thus the new cochain complex is simply the augmented cochain complex for the $\mathcal{A}$-peripheral complex (with indices shifted by one). Since the $\mathcal{A}$-peripheral complex is contractible (Lemma 4.4) its reduced cohomology is trivial hence this cochain complex is acyclic. □

5. Computing ranks via planted forests

We now know that the E_2 page consists of a single column and the entries are $\mathbb{Z}^i$ where i counts the number of essential sets of one-two basis elements. By its definition, every essential set of compatible one-two basis elements forms the basis of a hypertree with at most one edge of order > 2. These hypertrees can be described using planted forests.

Given a planted forest on $[n]$, f, let τ_f be the hypertree whose edges consist of the edges of f and one additional (hyper)edge consisting of the roots of f. That is, one forms the hypertree τ_f by gathering the roots of f. (See Figure 11.)

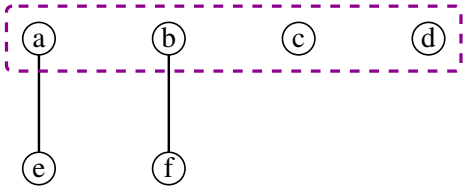

Figure 11. Gathering a planted forest with four components on [6] into a hypertree

Lemma 5.1. *A rank q hypertree τ is* essential *if and only if:*

1. *τ is formed by gathering the roots of a planted forest f on $[n]$ with (n-q) components; and*
2. *The vertex labelled 1 is a root of f and it is not the root of the tree containing the vertex labelled 2.*

There is a bijection between essential sets of one-two generators with q elements and rank q essential hypertrees. Thus it suffices to count the rank q essential hypertrees. The number of $(n-q)$ component planted forests on $[n]$ is

$$p_{n-q}(n) = \binom{n-1}{q} n^q$$

(See Proposition 5.3.2 of [19].)

Because of our second condition in Lemma 5.1—on the location of the vertex labelled by 1—not all hypertrees formed by gathering forests are essential. If any of the non-root vertices are labelled 1 we would be unhappy. Similarly, referring to Figure 11, if $a = 1$ then we would not want $e = 2$. That is, 2 should not label a vertex in the tree rooted by 1.

So given an unlabelled planted forest with k components and n vertices, what fraction of the labellings lead to essential hypertrees? First, place 2. There are then $(n-1)$ places you could put 1. Of these only the $(n-q-1)$-roots that are not the roots of the tree containing 2 will lead to essential hypertrees. Thus the number of essential hypertrees can be gotten by taking all hypertrees and multiplying by $\frac{n-q-1}{n-1}$.

Lemma 5.2. *The number of essential hypertrees in HT_n of rank q is*

$$\binom{n-2}{q} n^q \ .$$

Therefore $\mathrm{rk}\,[H^q(\mathrm{OP}\Sigma_n, \mathbb{Z})] = \binom{n-2}{q} n^q$.

When q is maximal (i.e. $q = n-2$) the hypertrees under consideration aer ordinary trees, all of them are essential, and Lemma 5.2 reproves the standard count of n^{n-2} for the number of trees with vertex set $[n]$.

Lemma 5.2 establishes our Main Theorem. In order to prove the Brownstein-Lee Conjecture we need a bit more information. It follows from McCool's presentation that $H_1(\mathrm{P}\Sigma_n, \mathbb{Z})$ is free abelian with one generator $[\alpha_{ij}]$ for each generator of $\mathrm{P}\Sigma_n$. The cohomology group $H^1(\mathrm{P}\Sigma_n, \mathbb{Z})$ is then generated by the dual basis $\{\alpha_{ij}^*\}$ where

$$\alpha_{ij}^*([\alpha_{kl}]) = \begin{cases} 1 & (k,l) = (i,j) \\ 0 & \text{otherwise} \end{cases} .$$

Similarly $\mathrm{OP}\Sigma_n$ is generated by $\{\bar{\alpha}_{ij} \mid i \neq j, i \neq 1, \text{ and } i \neq 2 \text{ if } j = 1\}$, its first homology is generated by the associated $[\bar{\alpha}_{ij}]$ and $H^1(\mathrm{OP}\Sigma_n, \mathbb{Z})$ is generated by the dual basis α_{ij}^*. The argument in Section 4, establishing that the E_2 page is concentrated in a single column (Proposition 4.1), gives an explicit description of the groups $E_2^{0,q}$. Namely, they are free abelian where the elements in our chosen generating set correspond to products of the elements $\alpha_{I,j}^* = \sum_{i \in I} \alpha_{ij}^*$. These cohomology classes $\alpha_{I,j}^*$ come from a generating set for the stablizer of an essential hypertree, hence it must be the case that for any $i \in I$, $i \neq 1$ and $i \neq 2$ if $j = 1$. Thus we get:

Proposition 5.3. *The algebra $H^*(\mathrm{OP}\Sigma_n, \mathbb{Z})$ is generated by the one-dimensional classes $\bar{\alpha}_{ij}^*$ where $i \neq j$, $i \neq 1$ and $i \neq 2$ if $j = 1$.*

6. Computing $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$

Starting with the short exact sequence

$$1 \to \mathbb{F}_n \to \mathrm{P}\Sigma_n \to \mathrm{OP}\Sigma_n \to 1$$

one can apply the Lyndon-Hochschild-Serre spectral sequence to compute the cohomology groups $H^i(\mathrm{P}\Sigma_n, M)$. Since $\mathrm{cd}(\mathbb{F}_n) = 1$ this spectral sequence is concentrated in the bottom two rows of the first quadrant (see Figure 12). In our case, it is easiest to understand the structure of this spectral sequence by thinking in terms of a fibration, and appealing to the Leray-Hirsch Theorem. There is a fiber bundle $p : B\mathrm{P}\Sigma_n \to B\mathrm{OP}\Sigma_n$ with fiber $B\mathbb{F}_n$ which gives a fibration (see Theorem 1.6.11 and Theorem 2.4.12 of [2].) The action of $\mathrm{P}\Sigma_n$ on $\mathbb{F}_n$ always sends a generator to a conjugate of itself and so the induced action on $H^1(\mathbb{F}_n, \mathbb{Z}) = \mathbb{Z}^n$ is trivial; therefore $E_2^{*,1} = H^*(\mathrm{OP}\Sigma_n, H^1(\mathbb{F}_n, \mathbb{Z})) = H^*(\mathrm{OP}\Sigma_n, \mathbb{Z}^n)$ with the action of $\mathrm{OP}\Sigma_n$ on $\mathbb{Z}^n$ being trivial. Hence the system of local coefficients in the spectral sequence is simple.

We remind the reader of the Leray-Hirsch Theorem, as presented in [10]:

Theorem 6.1 (Leray-Hirsch Theorem)**.** *Let $F \xrightarrow{\iota} E \xrightarrow{\rho} B$ be a fiber bundle such that*

1. *$H^n(F, \mathbb{Z})$ is a finitely generated free $\mathbb{Z}$-module for each n, and*
2. *There exist classes $c_j \in H^{k_j}(E, \mathbb{Z})$ whose restrictions $\iota^*(c_j)$ form a basis for $H^*(F, \mathbb{Z})$ in each fiber F.*

Then the map $\Phi : H^*(B, \mathbb{Z}) \otimes H^*(F, \mathbb{Z}) \to H^*(E, \mathbb{Z})$ given by

$$\sum_{i,j} b_i \otimes \iota^*(c_j) \mapsto \sum_{i,j} \rho^*(b_i) \cup c_j$$

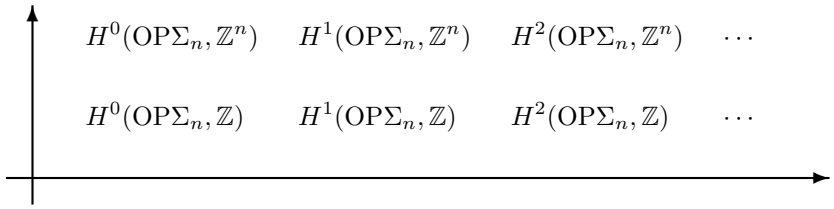

FIGURE 12. The E_2 page of the Lyndon-Hochschild-Serre spectral sequence for $1 \to \mathbb{F}_n \to \mathrm{P}\Sigma_n \to \mathrm{OP}\Sigma_n \to 1$.

is an isomorphism.

The first condition of the Leray-Hirsch Theorem is immediately satisfied since the kernel we are interested in is a free group. We turn then to establishing that the fiber $B\mathbb{F}_n$ is totally non-homologous to zero in $B\mathrm{P}\Sigma_n$ (with respect to $\mathbb{Z}$).

Lemma 6.2. *The space $B\mathbb{F}_n$ is totally non-homologous to zero in $B\mathrm{P}\Sigma_n$ with respect to $\mathbb{Z}$. In fact, let $c_{0,0}$ be a generator for $H^0(B\mathrm{P}\Sigma_n, \mathbb{Z})$ and let $c_{1,1}, \ldots, c_{1,n}$ denote the duals in $H^1(B\mathrm{P}\Sigma_n, \mathbb{Z}) = Hom((B\mathrm{P}\Sigma_n)_{ab}, \mathbb{Z})$ of maps corresponding to conjugating by the generators of $\mathbb{F}_n$. Then the collection $\{\iota^*(c_{r,s})\}$ forms an additive basis for $H^*(B\mathbb{F}_n, \mathbb{Z})$, where $\iota : B\mathbb{F}_n \to B\mathrm{P}\Sigma_n$ is inclusion.*

Proof. The inner automorphism of $\mathbb{F}_n$ given by conjugating by a basis element x_j is sent to the symmetric automorphism

$$\alpha_{[n]\setminus\{j\},j} = \alpha_{1j} \cdot \alpha_{2j} \cdots \hat{\alpha}_{jj} \cdots \alpha_{nj}$$

under the injection $\mathbb{F}_n \hookrightarrow \mathrm{P}\Sigma_n$. As the α_{ij} form a generating set for $\mathrm{P}\Sigma_n$ that projects to a minimal generating set for $H_1(\mathrm{P}\Sigma_n, \mathbb{Z})$, the map

$$(\mathbb{F}_n)_{ab} \to (\mathrm{P}\Sigma_n)_{ab}$$

is injective. Hence the map

$$H^1(\mathrm{P}\Sigma_n, \mathbb{Z}) = \mathrm{Hom}((\mathrm{P}\Sigma_n)_{ab}, \mathbb{Z}) \to H^1(\mathbb{F}_n, \mathbb{Z}) = \mathrm{Hom}((\mathbb{F}_n)_{ab}, \mathbb{Z})$$

is onto. Since the cohomology of $\mathbb{F}_n$ is only located in degrees 0 and 1, this means that the map $H^*(\mathrm{P}\Sigma_n, \mathbb{Z}) \to H^*(\mathbb{F}_n, \mathbb{Z})$ is onto. □

Having satisfied the hypotheses we may apply the Leray-Hirsch Theorem to obtain

Lemma 6.3. *The map*

$$H^*(\mathrm{OP}\Sigma_n, \mathbb{Z}) \otimes H^*(\mathbb{F}_n, \mathbb{Z}) \to H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$$

defined by

$$\sum_{j,s} b_j \otimes i^*(c_{r,s}) \mapsto \sum_{j,s} p^*(b_j) \cup c_{r,s}$$

is an isomorphism. Furthermore, the above spectral sequence has trivial differential d_2, and therefore the rank of $H^i(\mathrm{P}\Sigma_n, \mathbb{Z})$ is $\binom{n-1}{i} \cdot n^i$

Proof. That the map is an isomorphism follows from the Leray-Hirsch Theorem. To compute the ranks of the cohomology groups, we note that $H^i(\mathrm{P}\Sigma_n, \mathbb{Z}) = H^{i-1}(\mathrm{OP}\Sigma_n, \mathbb{Z}^n) \times H^i(\mathrm{OP}\Sigma_n, \mathbb{Z})$ which by the Main Theorem means

$$H^i(\mathrm{P}\Sigma_n, \mathbb{Z}) = \mathbb{Z}^{\binom{n-2}{i-1}\cdot n^{i-1}\cdot n} \times \mathbb{Z}^{\binom{n-2}{i}\cdot n^i} = \mathbb{Z}^{\left[\binom{n-2}{i-1}+\binom{n-2}{i}\right]\cdot n^i} = \mathbb{Z}^{\binom{n-1}{i}\cdot n^i}.$$

□

Thus at the level of abelian groups, we have the formula claimed in Theorem 6.7 below. At this point the only cause for caution is that the Leray-Hirsch Theorem does not immediately imply a *ring* isomorphism. This we establish via the next two results.

Corollary 6.4. *The integral cohomology of* $\mathrm{P}\Sigma_n$ *is generated by one-dimensional classes.*

Proof. Consider an $x = p^*(b) \cup c_{1,s} \in H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$. From 5.3, $H^*(\mathrm{OP}\Sigma_n, \mathbb{Z})$ is generated by one-dimensional classes and so there are $d_1, \ldots, d_t \in H^1(\mathrm{OP}\Sigma_n, \mathbb{Z})$ such that $b = d_1 \cup \cdots \cup d_t$. Hence $p^*(b) = p^*(d_1 \cup \cdots \cup d_t) = p^*(d_1) \cup \cdots \cup p^*(d_n)$ (cf. property (3.7) in Chapter V of [5]) is also a product of one-dimensional classes. So x is a product of one-dimensional classes. □

In their work on $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$, Brownstein and Lee established

Theorem 6.5 (Theorems 2.10 and 2.11 in [6])**.** *The relations*

1. $\alpha^*_{ij} \wedge \alpha^*_{ij} = 0$
2. $\alpha^*_{ij} \wedge \alpha^*_{ji} = 0$
3. $\alpha^*_{kj} \wedge \alpha^*_{ji} = (\alpha^*_{kj} - \alpha^*_{ij}) \wedge \alpha^*_{ki}$

hold among the one-dimensional classes α^*_{ij} *in* $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$*, and—in the case where* $n = 3$*—they present the algebra* $H^*(\mathrm{P}\Sigma_3, \mathbb{Z})$*.*

These relations give an upper bound on the rank of $H^i(\mathrm{P}\Sigma_n, \mathbb{Z})$.

Lemma 6.6. *The Brownstein-Lee relations imply that the rank of* $H^i(\mathrm{P}\Sigma_n, \mathbb{Z})$ *is at most* $\binom{n-1}{i} n^i$*.*

Proof. The first relations combined with the fact that one-dimensional classes generate $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$ imply that we may form a basis for $H^i(\mathrm{P}\Sigma_n, \mathbb{Z})$ using i-fold products of the α^*_{ij} with no repetitions. We may rewrite the third set of relations as

$$\alpha^*_{kj}\alpha^*_{ki} = \alpha^*_{kj}\alpha^*_{ji} + \alpha^*_{ij}\alpha^*_{ki}$$

implying that we never need to repeat the first index. An induction argument using the original expression for the third relation shows that any telescoping sequence $\alpha^*_{ij}\alpha^*_{jk}\alpha^*_{kl}\cdots\alpha^*_{st}$ is equivalent to a linear combination of terms, each of which includes α^*_{it}. But then applying the second relation we see that any cyclic product is trivial:

$$\alpha^*_{ij}\alpha^*_{jk}\alpha^*_{kl}\cdots\alpha^*_{st}\alpha^*_{ti} = \left(\sum \pm[\text{various } (i-1)\text{-fold products}]\alpha^*_{it}\right)\alpha^*_{ti} = 0.$$

Thus $H^i(\mathrm{P}\Sigma_n, \mathbb{Z})$ is generated as an abelian group by the i-fold products of the one-dimensional generators that do not repeat any α^*_{ij}, do not repeat a first index, and which do not contain any cyclic products.

We may encode such products as directed graphs on $[n]$ where including α^*_{jk} in the product adds an edge "$j \leftarrow k$". The fact that no cyclic products are allowed

implies the graph has no cycles, so it is a forest. The fact that no first index is repeated implies that each vertex is the target of at most one edge, hence there is a natural planting of the forest, with the roots corresponding to vertices that are not targets of any directed edge. Since there are i edges, and n vertices, the forest must have $(n-i)$ connected components. But the number of k-component forests on $[n]$ is $\binom{n-1}{k-1}n^{n-k}$ hence the number of $(n-i)$-component forests on $[n]$—which gives the upper bound on the rank of $H^i(\mathrm{P}\Sigma_n, \mathbb{Z})$—is

$$\binom{n-1}{(n-i)-1}n^{n-(n-i)} = \binom{n-1}{i}n^i.$$

□

We know that $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$ is a quotient of the algebra presented by Brownstein and Lee's relations, but we also know by Lemma 6.3 that the rank of $H^i(\mathrm{P}\Sigma_n, \mathbb{Z})$ is $\binom{n-1}{i}n^i$. Thus we have established the Brownstein-Lee Conjecture (the Corollary to our Main Theorem).

Theorem 6.7. *The cohomology of $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$ is generated by one-dimensional classes α^*_{ij} where $i \neq j$, subject to the relations*

1. $\alpha^*_{ij} \wedge \alpha^*_{ij} = 0$
2. $\alpha^*_{ij} \wedge \alpha^*_{ji} = 0$
3. $\alpha^*_{kj} \wedge \alpha^*_{ji} = (\alpha^*_{kj} - \alpha^*_{ij}) \wedge \alpha^*_{ki}$

In particular, the Poincaré series is $\mathfrak{p}(z) = (1+nz)^{n-1}$.

Example 6.8. The cohomology groups of $\mathrm{P}\Sigma_4$ are

$$H^i(\mathrm{P}\Sigma_4, \mathbb{Z}) = \begin{cases} \mathbb{Z}^{64} & i = 3 \\ \mathbb{Z}^{48} & i = 2 \\ \mathbb{Z}^{12} & i = 1 \\ \mathbb{Z} & i = 0 \end{cases}.$$

As was remarked in §2, viewed as a subgroup of $\mathrm{Aut}(\mathbb{F}_n)$, $\mathrm{P}\Sigma_n$ is contained in the subgroup IA_n. Magnus proved that IA_n is generated by the α_{ij}, which generate $\mathrm{P}\Sigma_n$, along with the automorphisms induced by

$$\theta_{ijk} = \begin{cases} x_i \to x_i[x_j, x_k] & \\ x_l \to x_l & l \neq i \end{cases}.$$

Aside from this generating set, little is known about IA_n. Fred Cohen and Jon Pakianathan, and independently Benson Farb, showed that $H_1(\mathrm{IA}_n)$ is free abelian of rank $n^2(n-1)/2$, and is generated by the classes $[\alpha_{ij}]$ and $[\theta_{ijk}]$. Krstić and McCool proved that IA_3 is not finitely presentable [12]. The Krstić-McCool result has recently been extended by Bestvina, Bux and Margalit to IA_n for $n \geq 3$; they also succeed in computing the cohomological dimension of IA_n and in showing that its top dimensional cohomology is not finitely generated [3]. Fred Cohen has pointed out to us that Theorem 6.7 gives some information about $H^*(\mathrm{IA}_n, \mathbb{Z})$:

Corollary 6.9. *The injection $\mathrm{P}\Sigma_n \hookrightarrow IA_n$ induces a split epimorphism*

$$H^*(IA_n, \mathbb{Z}) \twoheadrightarrow H^*(\mathrm{P}\Sigma_n, \mathbb{Z}).$$

Moreover, the suspension of $B\mathrm{P}\Sigma_n$ is homotopy equivalent to a bouquet of spheres, and it is a retract of the suspension of BIA_n.

Remark 3. The suspension of BP_n, where P_n is the pure braid group, is also homotopy equivalent to a wedge of spheres. (This is Corollary 3 of [18]; see also [1], where this property and its implications for K-theory are explored.) Like the inclusion $\mathrm{P}\Sigma_n \hookrightarrow \mathrm{IA}_n$, the inclusion $P_n \hookrightarrow \mathrm{P}\Sigma_n$ induces a surjection on cohomology: $H^*(\mathrm{P}\Sigma_n, \mathbb{Z}) \twoheadrightarrow H^*(P_n, \mathbb{Z})$, but this surjection does not split over $\mathbb{Z}$, but it does split over $\mathbb{Z}[\frac{1}{2}]$ (Proposition 4.3 of [6]).

Proof of Corollary 6.9, due to Fred Cohen. Consider the map $\mathrm{P}\Sigma_n \hookrightarrow \mathrm{IA}_n$ composed with the abelianization

$$\mathrm{IA}_n \to \bigoplus_{n\binom{n}{2}} \mathbb{Z} = H_1(\mathrm{IA}_n, \mathbb{Z}) \ .$$

The induced map from the first homology of $\mathrm{P}\Sigma_n$ to the first homology group of $\mathbb{Z}^{n\binom{n}{2}}$ is a split monomorphism. Thus the induced map on the level of integral cohomology is a split epimorphism, at least in degree 1. Since $H^1(\mathrm{P}\Sigma_n, \mathbb{Z})$ generates $H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$ as an algebra, it follows that

$$H^*(\bigoplus_{n\binom{n}{2}} \mathbb{Z}, \mathbb{Z}) \to H^*(\mathrm{P}\Sigma_n, \mathbb{Z})$$

is an epimorphism.

However, the cohomology of $\mathbb{Z}^{n\binom{n}{2}}$ is an exterior algebra on one-dimensional classes, some of which correspond to the generators for the cohomology of $\mathrm{P}\Sigma_n$. Furthermore, the single suspension of $B[\mathbb{Z}^{n\binom{n}{2}}]$ is a bouquet of spheres, with some subset of these corresponding exactly to the image of homology of $\mathrm{P}\Sigma_n$ (suitably reindexed). Denote this subset by $S_{\mathrm{P}\Sigma_n}$. We can then project from the suspension of $B[\mathbb{Z}^{n\binom{n}{2}}]$,

$$\Sigma(B[\mathbb{Z}^{n\binom{n}{2}}]) \to S_{\mathrm{P}\Sigma_n} \ .$$

The composition

$$\Sigma(B\mathrm{P}\Sigma_n) \to \Sigma(B[\mathbb{Z}^{n\binom{n}{2}}]) \to S_{\mathrm{P}\Sigma_n}$$

induces a homology isomorphism (of simply connected spaces) and thus is a homotopy equivalence. It follows that $\Sigma(B\mathrm{P}\Sigma_n)$ is a retract of $\Sigma(B\mathrm{IA}_n)$. □

Remark 4. Do the connections above give new information about the cohomology of IA_n? Using the Johnson homomorphism Alexandra Pettet has found a large number of cohomology classes in $H^2(\mathrm{IA}_n, \mathbb{Z})$ [17], all of which are in the image of the map $\wedge^2 H^1(\mathrm{IA}_n, \mathbb{Z}) \to H^2(\mathrm{IA}_n, \mathbb{Z})$. Pettet points out that as a consequence of Corollary 6.9 one gets a commutative diagram

$$\begin{array}{ccc} \wedge^2 H^1(\mathrm{IA}_n, \mathbb{Z}) & \longrightarrow & \wedge^2 H^1(\mathrm{P}\Sigma_n, \mathbb{Z}) \\ \downarrow & & \downarrow \\ H^2(\mathrm{IA}_n, \mathbb{Z}) & \longrightarrow & H^2(\mathrm{P}\Sigma_n, \mathbb{Z}) \end{array}$$

where all the maps, excluding the left edge, are surjections. Thus the classes arising from Corollary 6.9 are included in the classes Pettet has found.

Another curious point is that Corollary 6.9 gives a non-trivial map

$$H_*(BIA_n) \to T[\bar{H}_*(B\mathrm{P}\Sigma_n)]$$

where T[V] denotes the tensor algebra generated by V. This process also gives a map out of $B\mathrm{IA}_n$ to a highly non-trivial space, which *might* give more information about $B\mathrm{IA}_n$. The map is

$$B\mathrm{IA}_n \to \Omega\Sigma(B\mathrm{P}\Sigma_n)$$

factoring the Freudenthal suspension $B\mathrm{P}\Sigma_n \to \Omega\Sigma(B\mathrm{P}\Sigma_n)$. By arguments similar to those above, the composite

$$B\mathrm{P}\Sigma_n \to B\mathrm{IA}_n \to \Omega\Sigma(B\mathrm{P}\Sigma_n)$$

induces a split epimorphism in cohomology. For information on such arguments, in particular how they apply to the subgroup of $\mathrm{P}\Sigma_n$ generated by $\{\alpha_{ij} \mid i > j\}$, see [7].

References

[1] A. Adem, D. Cohen, and F. R. Cohen. On representations and K-theory of the braid groups. *Math. Ann.*, 326(3):515–542, 2003.

[2] D.J. Benson. *Representations and cohomology II: Cohomology of groups and modules*, volume 31 of *Cambridge studies in advanced mathematics*. Cambridge University Press, Cambridge, Great Britain, 1991.

[3] M. Bestvina, K-U. Bux, and D. Margalit. Dimension of the Torelli group for $\mathrm{Out}(\mathbb{F}_n)$. preprint.

[4] Noel Brady, Jon McCammond, John Meier, and Andy Miller. The pure symmetric automorphisms of a free group form a duality group. *J. Algebra*, 246(2):881–896, 2001.

[5] Kenneth S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.

[6] Alan Brownstein and Ronnie Lee. Cohomology of the group of motions of n strings in 3-space. In *Mapping class groups and moduli spaces of Riemann surfaces (Göttingen, 1991/Seattle, WA, 1991)*, volume 150 of *Contemp. Math.*, pages 51–61. Amer. Math. Soc., Providence, RI, 1993.

[7] F. R. Cohen and J. Pakianathan. On subgroups of the automorphism group of a free group and associated graded Lie algebras. In preparation.

[8] Donald J. Collins. Cohomological dimension and symmetric automorphisms of a free group. *Comment. Math. Helv.*, 64(1):44–61, 1989.

[9] Deborah L. Goldsmith. The theory of motion groups. *Michigan Math. J.*, 28(1):3–17, 1981.

[10] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[11] Craig Jensen, Jon McCammond, and John Meier. The Euler characteristic of the Whitehead automorphism group of a free product. To appear, *TAMS*.

[12] Sava Krstić and James McCool. The non-finite presentability of $\mathrm{IA}(F_3)$ and $\mathrm{GL}_2(\mathbf{Z}[t, t^{-1}])$. *Invent. Math.*, 129(3):595–606, 1997.

[13] Jon McCammond and John Meier. The hypertree poset and the l^2-Betti numbers of the motion group of the trivial link. *Math. Ann.*, 328(4):633–652, 2004.

[14] J. McCool. On basis-conjugating automorphisms of free groups. *Canad. J. Math.*, 38(6):1525–1529, 1986.

[15] Darryl McCullough and Andy Miller. Symmetric automorphisms of free products. *Mem. Amer. Math. Soc.*, 122(582):viii+97, 1996.

[16] Alexandra Pettet. Finiteness properties for the kernel of pure motions of n unlinked loops.

[17] Alexandra Pettet. The Johnson homomorphism and the second cohomology of IA_n. *Algebr. Geom. Topol.*, 5:725–740 (electronic), 2005.

[18] Ch. Schaper. Suspensions of affine arrangements. *Math. Ann.*, 309(3):463–473, 1997.

[19] Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999.

Dept. of Mathematics, University of New Orleans, New Orleans, LA 70148
E-mail address: jensen@math.uno.edu

Dept. of Mathematics, University of California, Santa Barbara, CA 93106
E-mail address: jon.mccammond@math.ucsb.edu

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