

THE HYPERTREE POSET AND THE ℓ^2 -BETTI NUMBERS OF THE MOTION GROUP OF THE TRIVIAL LINK

JON MCCAMMOND¹ AND JOHN MEIER²

ABSTRACT. We give explicit formulae for the Euler characteristic and ℓ^2 -cohomology of the group of motions of the trivial link, or isomorphically the group of free group automorphisms that send each standard generator to a conjugate of itself. The method is primarily combinatorial and ultimately relies on a computation of the Möbius function for the poset of labelled hypertrees.

1. INTRODUCTION

Classic combinatorial group theory, such as what is described in [11] or [12], uses relatively elementary combinatorics to study infinite groups. For instance, small cancellation theory is the study of groups with finite presentations whose associated Whitehead graph has large girth. Increasingly there is a need to use more sophisticated combinatorial arguments to establish topological properties of infinite groups. Here we compute the Möbius function of a poset of labelled hypertrees in order to explicitly describe the ℓ^2 -Betti numbers of the motion group of a trivial n -component link. In earlier work a recursive atom ordering was used to compute these groups' cohomology with group ring coefficients [1]. While it is a speculative claim at this point, these two examples indicate that there may be general applicability of enumerative combinatorics in the study of group cohomology.

The poset we study consists of labelled hypertrees. Let $[n] = \{1, \dots, n\}$ and let HT_n be the set of all hypertrees with vertices labelled by $[n]$. The elements of HT_n admit a partial ordering where $\tau \leq \tau'$ if the hyperedges of τ' are contained in the hyperedges of τ . This poset has an element $\hat{0}$ that is less than all other elements; if $\hat{0}$ is removed we denote the resulting poset as $\text{HT}_n^{[\infty]}$; if a formal $\hat{1}$ is added to HT_n we denote the resulting poset as $\widehat{\text{HT}}_n$. Full definitions and background information on this poset can be found in §2. Our main combinatorial result is the following computation of the Möbius function of this poset. (The result stated below, combined with various lemmas in §2, imply a complete computation.)

Theorem 1.1. *For every $n \geq 0$,*

$$\mu_{\widehat{\text{HT}}_{n+1}}(\hat{0}, \hat{1}) = \tilde{\chi}(\text{HT}_{n+1}^{[\infty]}) = (-1)^n n^{n-1} .$$

The group we consider can be viewed as a higher dimensional version of the braid group. Let L_n denote n unknotted, unlinked circles in S^3 . Much as one can describe the n -strand braid group as the motion group of n points in a disk, one can define the motion group Σ_n of L_n in S^3 . Roughly speaking, an element of Σ_n

Date: December 2, 2003.

¹Partially supported by NSF grant no. DMS-0101506

²Partially supported by an AMS Centennial Research Fellowship

consists of an isotopy of S^3 that takes L_n back to L_n , with the circles possibly being permuted. In other words, elements of Σ_n are finite length movies consisting of n circles moving about in space (see §6 for a formal definition). The group operation is concatenation and inverses are formed by running the movie backward. The index $n!$ subgroup of motions where the n components of L_n return to their original positions is the *pure* motion group, denoted $P\Sigma_n$. For example, Figure 1 indicates two elements of $\Sigma_3 \setminus P\Sigma_3$. The first is gotten by moving from the top configuration to the bottom as indicated by the arrows; the second corresponds to moving from the bottom to the top along the righthand side. The product of these two elements, visualized by the entire circuit, is in $P\Sigma_3$.




FIGURE 1. An element of Σ_3 and $P\Sigma_3$

The group $P\Sigma_n$ admits an action on a contractible complex MM_n whose fundamental domain is the geometric realization of HT_n . (This is described in §6.) The action of $P\Sigma_n$ on MM_n is not free, or even proper, so one cannot immediately move from topological properties of MM_n to properties of the group $P\Sigma_n$. However, it is possible to gain non-trivial cohomological information from this action. In particular, we use the calculation of the Möbius function of \widehat{HT}_n to determine the ℓ^2 -Betti numbers of $P\Sigma_n$. Recalling Atiyah's result that the ordinary Euler characteristic is equivalent to the ℓ^2 -Euler characteristic for all groups of finite type [10, Theorem 1.35.2], it follows that this computation yields the ordinary Euler characteristic of $P\Sigma_n$, which was not previously known.

Theorem 1.2. *The reduced ℓ^2 -cohomology groups of $P\Sigma_{n+1}$ are*

$$\mathcal{H}^i(P\Sigma_{n+1}) = \begin{cases} \mathbb{Z}^{n^n} \otimes \ell^2(P\Sigma_{n+1}) & \text{if } i = n \\ 0 & \text{otherwise} \end{cases}$$

and therefore the ℓ^2 -Betti numbers of $P\Sigma_{n+1}$ are all trivial except in top dimension, where

$$\chi(P\Sigma_{n+1}) = (-1)^n b_n^{(2)} = (-1)^n n^n .$$

The ℓ^2 -Betti numbers can be extended to non-free group actions (cf §6.5 in [10]) and these ℓ^2 -Betti numbers are multiplicative (cf Theorem 6.54 in [10]) hence we may immediately extend the result of Theorem 1.2 to the over group Σ_{n+1} .

Corollary 1.3. *The ℓ^2 -Betti numbers of Σ_{n+1} are trivial except in top dimension where*

$$\chi(\Sigma_{n+1}) = (-1)^n b_n^{(2)} = (-1)^n \frac{n^n}{(n+1)!} .$$

The study of ℓ^2 -cohomology is relatively recent and there are few groups where such concrete information is available. In fact, most results concentrate on vanishing theorems for the ℓ^2 -Betti numbers. One notable computation is due to Davis and Leary who computed the ℓ^2 -cohomology of the right angled Artin groups (and conjecturally all Artin groups) [5]. Our approach is modelled on a spectral sequence argument they outline in their final section.

The fact that the ℓ^2 -cohomology of Σ_{n+1} is non-trivial contrasts with the fact that the ℓ^2 -cohomology of the braid groups is trivial [5]. The fact that it is concentrated in top dimension contrasts the fact that for closed hyperbolic manifolds the ℓ^2 -Betti numbers are trivial except in middle dimension (cf Theorem 1.62 in [10]) while it is consonant with the recent result that for “very thick” buildings the ℓ^2 -cohomology is concentrated in top dimension [6].

The structure of this paper is as follows. In §§2–5 we describe the poset HT_n , review Möbius functions, and compute the Möbius function of $\widehat{\text{HT}}_n$. In §6 we give the full definition of $P\Sigma_n$ and describe the connection between it and the poset HT_n . In §7 we quickly review the relevant facts about ℓ^2 -cohomology, which we then use in §8 to establish the ℓ^2 -Betti number formulae.

We thank Mike Davis, Ian Leary, Thomas Schick and Derek Smith for their helpful comments during the development of this paper.

2. HYPERTREES

The notions of hypergraphs and hypertrees are natural generalizations of graphs and trees, respectively. In this section we review the basic definitions and define a partial ordering that turns the set of hypertrees on $[n]$ into a poset.

Definition 2.1 (Hypergraphs and hypertrees). A *hypergraph* Γ is an ordered pair (V, E) where V is the set of vertices and E is a collection of subsets of V each containing at least two elements. An edge $e \in E$ that contains i vertices is called an *i -edge*, and the number of edges which contain a vertex v is called its *degree*. When it is necessary to highlight the fact that Γ is not necessarily a graph, the elements of E are sometimes called *hyperedges*. A *walk* in a hypergraph Γ is a sequence $v_0, e_1, v_1, \dots, v_{n-1}, e_n, v_n$ where for all i , $v_i \in V$, $e_i \in E$ and for each e_i , $\{v_{i-1}, v_i\} \subset e_i$. If for every pair of vertices v and v' there is a walk in Γ starting

at v and ending at v' , then Γ is called *connected*. A walk is a *cycle* if it 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 connected hypergraph with no cycles is called a *hypertree*. Notice that the no cycle condition implies that distinct edges in Γ have at most one vertex in common.

Remark 2.2 (Number of hypertrees). The number of elements in HT_n had been calculated both by Kalikow [9] and by Smith and Warne [17], and it grows quite rapidly. For example, the number of hypertrees on the set $[n]$ for $n = 2, \dots, 9$ are

$$1, 4, 29, 311, 4447, 79745, 1722681, 43578820 .$$

The general formula is $|\text{HT}_n| = \sum_k n^{k-1} S(n-1, k)$, where $S(n, k)$ denotes the Stirling numbers of second kind. See also the remark after Corollary 3.4.

The partial ordering on HT_n closely resembles the partial ordering of the partition lattice.

Definition 2.3 (Hypertree poset). Let τ and τ' be two hypertrees on the set $[n]$. We say that $\tau \leq \tau'$ if each edge of τ' is contained in an edge of τ . We write $\tau < \tau'$ if $\tau \leq \tau'$ but $\tau \neq \tau'$. In Figure 4, $A < B$, $B < C$, $B < D$ while C and D are incomparable. The set of all hypertrees on $[n]$ forms a poset which we denote HT_n . The poset HT_n is a graded poset and the hypertrees at height i are precisely those hypertrees with $i+1$ edges. This number is called the *rank* of τ in [13] and [1]. Notice that HT_n contains a unique minimal element (with height 0) that has only a single edge containing all of $[n]$. This is the *nuclear vertex* of [13] and we denote it by $\hat{0}$. At the other extreme HT_n has as many maximal elements as there are trees with vertex set labelled by $[n]$, that is, there are n^{n-2} maximal elements.

One can restrict this poset by rank selection, in particular, by removing the minimal element. Denote $\text{HT}_n \setminus \hat{0}$ by $\text{HT}_n^{[\infty]}$. Lemma 6.5 explains why this notation is not as unnatural as it may first appear.

One can extend HT_n by adding a formal element $\hat{1}$, which lies, by definition, above every element in HT_n . Denote the resulting bounded poset by $\widehat{\text{HT}}_n$.

Example 2.4 (HT_4). The poset HT_4 is illustrated in Figure 2. The letters used to label an element of HT_4 in Figure 2 are meant to indicate which of the trees in Figure 4 this element resembles once the vertex labels are ignored. Thus the four vertices labelled D represent trees isomorphic to the hypertree labelled D in Figure 4 and the vertex labelled A , the nuclear vertex for HT_4 , represents unique tree in HT_4 with a single edge of size 4. Using this convention, all of the unlabelled vertices in the top row should be labelled C and all of the unlabelled vertices in the middle row should be labelled B . For another description of HT_4 see Figure 8 in [13], and for a detailed description of HT_5 see [1].

The structure of the posets $\widehat{\text{HT}}_n$ is well-suited to a recursive analysis because of the following result.

Lemma 2.5 (Intervals). *Let τ be a hypertree on the set $[n]$.*

- (1) *The half-open interval $[\tau, \hat{1})$ is a direct product of hypertree posets, with one factor of HT_j for each edge in τ with size j .*
- (2) *The interval $[\hat{0}, \tau]$ is a direct product of partition lattices, with one factor Π_j for each vertex in τ with degree j .*




FIGURE 2. The poset HT_4

Proof. The first result is immediate once it is realized that the hypertrees above τ are obtained by replacing each hyperedge e with another hypertree on the set e , and that these replacements are independent of each other. Similarly, the hypertrees below τ in the ordering are obtained by merging the edges containing a common vertex v according to some partition and these mergings can also be carried out independently. \square

Although we will not need it, we record the following general classification result.

Corollary 2.6. *Let $\tau < \tau'$ be elements of $\widehat{\text{HT}}_n$. When $\tau' \neq \hat{1}$, the interval $[\tau, \tau']$ is a direct product of partition lattices, and when $\tau' = \hat{1}$ the interval $[\tau, \tau']$ is a direct product of hypertree posets with a new maximum element added to the result.*

Definition 2.7 (Edge sizes). If τ is a hypertree on the set $[n]$, let $\text{EDGESIZES}(\tau)$ denote the multiset which records the sizes of the edges in τ . For example, the hypertrees shown in Figure 4 have $\text{EDGESIZES}(A) = \{4\}$, $\text{EDGESIZES}(B) = \{2, 3\}$, and $\text{EDGESIZES}(C) = \text{EDGESIZES}(D) = \{2, 2, 2\}$.

Using this notation, the first part of Lemma 2.5 can be rewritten as follows. For each $\tau \in \text{HT}_n$,

$$[\tau, \hat{1}] = \prod_{i \in \text{EDGESIZES}(\tau)} \text{HT}_i$$

We conclude this section by summarizing in Theorem 2.9 the main properties of the hypertree posets that have previously been established. We first recall the following definition.

Definition 2.8 (Cohen-Macaulay). A poset \mathcal{P} is *Cohen-Macaulay* if its geometric realization $|\mathcal{P}|$ is Cohen-Macaulay. That is, for any closed simplex $\sigma \subset |\mathcal{P}|$ (including the empty simplex)

$$\tilde{H}_i(\text{Lk}(\sigma)) = \begin{cases} 0 & i \neq \dim(|\mathcal{P}|) - \dim(\sigma) - 1 \\ \text{torsion free} & i = \dim(|\mathcal{P}|) - \dim(\sigma) - 1 \end{cases}$$

where the empty simplex has formal dimension $= -1$.

Theorem 2.9. *For each $n \geq 1$, the poset $\widehat{\text{HT}}_n$ is a finite, bounded, graded lattice that is Cohen-Macaulay.*

Proof. The posets $\widehat{\text{HT}}_n$ are easily seen to be finite and bounded, the grading by ranks was first established by McCullough and Miller in [13]. That $\widehat{\text{HT}}_n$ is Cohen-Macaulay was established in [1] by showing that $\widehat{\text{HT}}_n$ admits a recursive atom

ordering, hence it is shellable, hence Cohen-Macaulay. Finally, the lattice property was essentially established in [13]. Strictly speaking, McCullough and Miller only showed that any subset of HT_n which has an upperbound, also has a least upper bound. This shows that joins in $\widehat{\text{HT}}_n$ are well-defined, but finite bounded posets with well-defined joins or meets are always lattices [15, Proposition 3.3.1]. Viewing the underlying objects as hypertrees — a different perspective than in [13] — makes this argument more transparent. It is easy to see that the meet of two hypertrees τ and τ' is well-defined by taking the union of the two sets of edges and then combining edges when they overlap too much or when a series of them form a cycle. The hypertree which results is clearly a lower bound for τ and τ' and just as clearly a greatest lower bound since all of these operations are forced. \square

As was commented in the proof above, McCullough and Miller did not view HT_n as a poset of hypertrees, but rather the results in Theorem 2.9 were originally established for the poset of $[n]$ -labelled bipartite trees [13]. In order to aid someone working through the original literature we quickly describe how to convert between the poset of hypertrees on $[n]$ and the poset of $[n]$ -labelled bipartite trees.

Definition 2.10 ($[n]$ -labelled bipartite trees). An $[n]$ -labelled bipartite tree is a tree T together with a bijection from $[n]$ to a subset of its vertex set such that the image of $[n]$ includes all of the vertices of valence 1 and for every edge in T exactly one of its endpoints lies in the image of $[n]$. The vertices that lie in the image of $[n]$ are called *labelled vertices* and the others are called *unlabelled vertices*. Two labelled bipartite trees are considered to be equivalent if there is a label preserving graph isomorphism between them. Several $[4]$ -labelled bipartite trees are shown in Figure 3.




FIGURE 3. Examples of $[4]$ -labelled bipartite trees.

The correspondence between $[n]$ -labelled bipartite trees and hypertrees on $[n]$ comes from identifying unlabelled vertices in the bipartite trees with hyperedges in the hypertrees. The hyperedge corresponding to an unlabelled vertex u is the subset of $[n]$ consisting of labels of the vertices connected to u . For example, in Figure 3, the $[4]$ -labelled bipartite tree labelled B is associated to the hypertree two edges, $\{1, 2, 3\}$ and $\{3, 4\}$. Conversely, a hypertree τ can be converted to an $[n]$ -labelled bipartite tree by starting with vertices labelled by $[n]$ and then coning off the vertices in each edge e in τ ; the cone point is then an unlabelled vertex. The hypertrees corresponding to the $[4]$ -labelled bipartite trees in Figure 3 are shown in Figure 4.




FIGURE 4. Examples of hypertrees on $[4]$.

McCullough and Miller define a partial ordering on $[n]$ -labelled trees via the following process. Let τ be an $[n]$ -labelled bipartite tree with distinct edges e_1 and e_2 sharing a common labelled endpoint v and with unlabelled endpoints u_1 and u_2 respectively. If τ' is the tree obtained from τ by identifying the edges e_1 and e_2 as well as the vertices u_1 and u_2 , then we say that the tree τ' is obtained from τ by *folding at v* . The partial order is the transitive closure of this operation. The collection of all $[n]$ -labelled bipartite trees with this partial ordering is called the *Whitehead poset*, denoted W_n [13]. To see that the Whitehead poset W_n is isomorphic to the hypertree poset HT_n , one simply notes that folding together unlabelled vertices u_1 and u_2 corresponds to replacing the hyperedges corresponding to u_1 and u_2 with a single hyperedge containing the union of their vertices. Thus we have established:

Lemma 2.11. *The Whitehead poset W_n is isomorphic to the hypertree poset HT_n .*

3. EXPONENTIAL GENERATING FUNCTIONS

Associated to the posets HT_n are two multivariable generating functions, T and R , whose properties capture the recursive nature of these posets. Throughout this section we follow the notational conventions established in [9].

Definition 3.1 (Hypertree generating function). Let τ be a hypertree on the set $[n]$ and let λ_i denote the number of i -edges in τ . The *weight* of τ will be the multivariate monomial $\text{WEIGHT}(\tau) = u_2^{\lambda_2} u_3^{\lambda_3} \cdots u_n^{\lambda_n}$. In other words,

$$\text{WEIGHT}(\tau) = \prod_{i \in \text{EDGE SIZES}(\tau)} u_i.$$

Next, let T_n denote the polynomial which is the sum of $\text{WEIGHT}(\tau)$ for all $\tau \in \text{HT}_n$, and let T denote the exponential generating function for the polynomials T_n . In other words T is a formal power series depending on variables t and u_j , $j \geq 1$ and

$$T = \sum_{n \geq 1} T_n \frac{t^n}{n!}.$$

For example, the polynomial for T_4 is $u_4 + 12u_2u_3 + 16u_2^3$ since there is only one hypertree with a single 4-edge (the one of type *A*), twelve hypertrees with one 2-edge and one 3-edge (those of type *B*) and sixteen hypertrees with three 2-edges (twelve of type *C* and four of type *D*). We have listed some values of T_n in Table 1.

n	T_n
1	1
2	u_2
3	$u_3 + 3u_2^2$
4	$u_4 + 12u_2u_3 + 16u_2^3$
5	$u_5 + 20u_2u_4 + 15u_3^2 + 150u_2^2u_3 + 125u_2^4$

TABLE 1. T_n for small values of n

Definition 3.2 (Rooted hypertree generating function). There are analogous polynomials R_n which are obtained by summing the weights of all *rooted* hypertrees, and there is an obvious relation between T_n and R_n , namely, $R_n = n \cdot T_n$. When the polynomials R_n are collected into an exponential generating function R , as above, then $R = t \frac{\partial T}{\partial t}$. The values of R_n for small n have been recorded in Table 2.

n	R_n
1	1
2	$2u_2$
3	$3u_3 + 9u_2^2$
4	$4u_4 + 48u_2u_3 + 64u_2^3$
5	$5u_5 + 120u_2u_4 + 75u_3^2 + 750u_2^2u_3 + 625u_2^4$

TABLE 2. R_n for small values of n

The main result we need about R is the functional equation established by Louis Kalikow in [9].

Theorem 3.3 (Kalikow). *With R and u_j defined as above, the following equation holds in the ring of formal power series with commuting variables t , and u_j , $j \geq 1$.*

$$R = te^y \text{ where } y = \sum_{j \geq 1} u_{j+1} \frac{R^j}{j!}.$$

When each variable u_j is replaced with 1, the function R reduces to a formal power series in t alone. Call this $R(t)$ and notice that the coefficient of $\frac{t^n}{n!}$ in $R(t)$ counts the number of rooted hypertrees on the set $[n]$. In particular, Theorem 3.3 then reduces to the earlier result of W. D. Smith [17, Theorem 3.14].

Corollary 3.4 (W. D. Smith). *The formal power series $R(t)$ satisfies the equation*

$$R(t) = te^{e^{R(t)} - 1}.$$

One method of deriving the Stirling number formula given in Remark 2.2 is to apply Lagrange inversion to the equation in Corollary 3.4.

4. MÖBIUS FUNCTIONS

In this section we use Möbius functions to establish a second, particularly nice, specialization of Theorem 3.3. We begin with a quick review of zeta and Möbius functions on posets.

Definition 4.1 (Zeta and Möbius functions). Let P be a finite poset. The *zeta function* for P is a function $\zeta_P : P \times P \rightarrow \mathbb{C}$ where $\zeta(x, y)$ is 1 if $x \leq y$ in P and 0 otherwise. If the elements of P are linearly ordered in a manner consistent with the poset ordering, then ζ can be represented as an upper triangular matrix with 1s down the diagonal and 1s above the diagonal indicating the poset structure. The *Möbius function* of P is the function $\mu_P : P \times P \rightarrow \mathbb{C}$ represented by the inverse of this matrix. Alternatively, it can be described as the unique function with the properties (1) $\mu_P(x, x) = 1$ for all $x \in P$, (2) $\mu_P(x, y) = 0$ if $x \not\leq y$ and

$$(3) \quad \sum_{z \in [x, y]} \mu_P(x, z) = 0 \text{ for all } x < y ;$$

dually, the Möbius function also has the property

$$(3') \quad \sum_{z \in [x, y]} \mu_P(z, y) = 0 \text{ for all } x < y .$$

In order to complete our ℓ^2 -cohomology calculation, the most important property of Möbius functions is

Lemma 4.2 (Proposition 3.8.6 in [15]). *Let P be a finite poset. Then*

$$\mu_{\hat{P}}(\hat{0}, \hat{1}) = \tilde{\chi}(|P|) .$$

In order to carry complete our Möbius function computations, we need to relate the Möbius function of the interval $[\tau, \hat{1}]$ to the Möbius function of the posets into which it can be factored (Lemma 2.5). For this, we use the idea of a sum function.

Definition 4.3 (Sum function). If P is a finite poset with a unique minimum element, then we define a *sum function*, $s(P) = \sum_{x \in P} \mu(\hat{0}, x)$ that simply adds up all of the Möbius values which begin at the minimum element. Notice that if \hat{P} denotes the poset derived from P by the addition of a new maximum element, $\hat{1}$, then $\mu_{\hat{P}}(\hat{0}, \hat{1}) = -s(P)$ by the third property of Möbius functions given in Definition 4.1.

The key property of sum functions is that they behave well with the taking of direct products.

Lemma 4.4. *If $P_i, i \in [k]$ is a list of finite posets each with a unique minimal element and $Q = \prod_{i=1}^k P_i$, then $s(Q) = \prod_{i=1}^k s(P_i)$.*

Proof. The proof is a standard manipulation of Möbius functions. In the equations below we use the fact that Möbius function between two elements in a direct product is a direct product of the Möbius functions taken coordinate by coordinate. See [15, Proposition 3.8.2].

$$\begin{aligned} s(Q) &= \sum_{(x_1, x_2, \dots, x_k) \in Q} \mu_Q(\hat{0}, (x_1, x_2, \dots, x_k)) \\ &= \sum_{(x_1, x_2, \dots, x_k) \in Q} \prod_{i=1}^k \mu_{P_i}(\hat{0}, x_i) \\ &= \prod_{i=1}^k \left(\sum_{x_i \in P_i} \mu_{P_i}(\hat{0}, x_i) \right) \\ &= \prod_{i=1}^k s(P_i) \end{aligned}$$

□

Definition 4.5 (Möbius numbers). For each hypertree τ on the set $[n]$ let μ_τ be an abbreviation for $\mu_{\widehat{\text{HT}}_n}(\tau, \hat{1})$ and let m_τ be its absolute value. In the special case where τ is the nuclear vertex in HT_n , we abbreviate these numbers as μ_n and m_n . Similarly, we abbreviate $s(P)$ where $P = [\tau, \hat{1}]$ as s_τ and $s(\text{HT}_n)$ as s_n .

Using Lemma 4.4 we can rewrite μ_τ in terms of the variables m_j .

Lemma 4.6. *If τ is a hypertree on the set $[n]$ with r edges, then*

$$\mu_\tau = (-1)^n \prod_{i \in \text{EDGE SIZES}(\tau)} (-m_i) .$$

Proof. By the remark in Definition 4.1 and Lemma 4.4

$$-\mu_\tau = s_\tau = \prod_{i \in \text{EDGE SIZES}(\tau)} s_i = (-1)^r \prod_{i \in \text{EDGE SIZES}(\tau)} \mu_i .$$

Thus μ_τ and $\prod_{i \in \text{EDGE SIZES}(\tau)} m_i$ at least agree up to a sign. To analyze the sign of μ_τ we first note that the sign of μ_i is $(-1)^{i-1}$ because HT_i is Cohen-Macaulay and Möbius functions of Cohen-Macaulay posets alternate in sign [15, Proposition 3.8.11], or more directly, because the reduced Euler characteristic of a Cohen-Macaulay poset is plus or minus the rank of its top dimensional homology, where the sign is determined by the dimension. Thus $\mu_i = (-1)^{i-1} m_i$. Next, an easy induction shows that $\sum_{i \in \text{EDGE SIZES}(\tau)} (i-1) = n-1$. Putting this all together we have

$$\begin{aligned} -\mu_\tau &= (-1)^r \prod_{i \in \text{EDGE SIZES}(\tau)} \mu_i \\ &= (-1)^r \prod_{i \in \text{EDGE SIZES}(\tau)} (-1)^{i-1} m_i \\ &= (-1)^{r+n-1} \prod_{i \in \text{EDGE SIZES}(\tau)} m_i \\ &= (-1)^{n-1} \prod_{i \in \text{EDGE SIZES}(\tau)} (-m_i) \end{aligned}$$

which completes the proof. \square

The following Corollary is now immediate.

Corollary 4.7. *If τ is a hypertree on $[n]$, then the weight of τ with each variable u_j replaced by $-m_j$ is the same as the value of $(-1)^n \mu_\tau$.*

Using this we can show that the specialization of T_n at the numbers $-m_j$ is particularly simple.

Lemma 4.8. *If T'_n denotes the number which results when for each $j \geq 1$ the variable u_j in T_n is replaced with the number $-m_j$, then $T'_n = (-1)^{n+1}$.*

Proof. First note that the sum of μ_τ for all $\tau \in \widehat{\text{HT}}_n$ is 0 by one of the defining properties of Möbius functions, so that the sum of μ_τ for $\tau \in \text{HT}_n$ is $-\mu_{\widehat{\text{HT}}_n}(\hat{1}, \hat{1}) = -1$. Combining this fact with Corollary 4.7 gives

$$T'_n = \sum_{\tau \in \text{HT}_n} \left(\prod_{i \in \text{EDGE SIZES}(\tau)} (-m_i) \right) = (-1)^n \sum_{\tau \in \text{HT}_n} \mu_\tau = (-1)^{n+1} .$$

\square

Corollary 4.9. *If $T'(t)$ and $R'(t)$ denote the formal power series in t alone which result when for each $j \geq 1$, the variable u_j is replaced with the number $-m_j$, then $T'(t) = 1 - e^{-t}$ and $R'(t) = te^{-t}$.*

Proof. By Lemma 4.8 and the definition of T , $T'(t) = \sum_{n \geq 1} (-1)^{n+1} \frac{t^n}{n!}$, which is the formal power series for $1 - e^{-t}$. The evaluation $R'(t) = te^{-t}$ is then calculated using the relation $R = t \frac{\partial T}{\partial t}$. \square

5. THE MAIN COMBINATORIAL ARGUMENT

We are now ready to establish our main combinatorial result.

Theorem 5.1 (Möbius function computation). *For every $n \geq 0$,*

$$\mu_{\widehat{\text{HT}}_{n+1}}(\hat{0}, \hat{1}) = (-1)^n n^{n-1}$$

Proof. Let $T'(t)$, $R'(t)$ and $y'(t)$ denote the evaluations of T , R , and y where for each $j \geq 1$, the variable u_j has been replaced with the number $-m_j$. By Theorem 3.3, $R'(t) = te^{y'(t)}$, and by Corollary 4.9, $R'(t) = te^{-t}$. Thus t and $-y'(t)$ must be equal in the ring of formal power series. If we start with the formula which defines $y'(t)$ and we then replace each occurrence of $R'(t)$ in this formula with te^{-t} we find that the following equation must also hold.

$$(1) \quad t = \sum_{j \geq 1} m_{j+1} \frac{(te^{-t})^j}{j!}$$

Writing out the formal power series on the right and equating coefficients, one sees that there is a unique solution to this equation since the variable m_j first occurs in the coefficient of t^{j-1} . In particular, we can write each m_j in terms of the previous m_i 's and solve for them sequentially. Finally, the formula

$$(2) \quad t = \sum_{j \geq 1} j^{j-1} \frac{(te^{-t})^j}{j!}$$

was established in [16, Equation (5.44), p.28] in the context of studying the recursive structures of rooted trees. Comparing Equations 1 and 2, we conclude $m_{j+1} = j^{j-1}$ for all $j \geq 1$ is the unique solution. As was remarked in the proof of Lemma 4.6, the sign of the Möbius function is easy to determine in any Cohen-Macaulay poset, whence the formula. \square

As a final remark, we note that Theorem 5.1 can be combined with Lemma 2.5 and the known Möbius functions for the partition lattices to determine the value of $\mu_{\widehat{\text{HT}}_n}(\tau, \tau')$ for arbitrary τ and τ' in $\widehat{\text{HT}}_n$.

6. THE GROUPS Σ_{n+1} , $\text{P}\Sigma_{n+1}$, AND $\text{OP}\Sigma_{n+1}$

We now return to the motion groups under consideration. One needs to be cautious in defining motion groups, so we recall the following definitions from [7]. Here L_{n+1} denotes the unlinked, unknotted $(n+1)$ -component link and S^3 is the 3-sphere.

- (1) $H(S^3)$ is the space of self-homeomorphisms of the 3-sphere (compact-open topology).
- (2) $H(S^3, L_{n+1})$ is the subspace of homeomorphisms ϕ with $\phi(L_{n+1}) = L_{n+1}$ (with orientations preserved), for some fixed embedding of $L_{n+1} \hookrightarrow S^3$.
- (3) A *motion* of L_{n+1} in S^3 is a path $\mu : [0, 1] \rightarrow H(S^3)$ with $\mu(0)$ being the identity and $\mu(1) \in H(S^3, L_{n+1})$.

- (4) Two motions μ and ν are *equivalent* if $\mu^{-1}\nu$ is homotopic to a stationary motion, that is, a motion contained in $H(S^3, L_{n+1})$.

The collection of motions of L_{n+1} in S^3 under the equivalence relation forms a group denoted Σ_{n+1} . If one colors the $(n+1)$ components of L_{n+1} then the color preserving motions forms the *pure* motion group $P\Sigma_{n+1}$.

Collins proved that the cohomological dimension of $P\Sigma_{n+1}$ is n [3]; Gutiérrez and Krstić have shown that $P\Sigma_{n+1}$ has a regular language of normal forms [8].

The group Σ_{n+1} can also be presented as a group of free group automorphisms. Let F_{n+1} be a free group with fixed basis $X = \{x_1, \dots, x_{n+1}\}$. The *symmetric automorphism group* of F_{n+1} , isomorphic to Σ_{n+1} , consists of those automorphisms that send each $x_i \in X$ to a conjugate of some $x_j \in X$. The *pure symmetric automorphism group* ($P\Sigma_{n+1}$) is the index $(n+1)!$ subgroup of Σ_{n+1} of symmetric automorphisms that send each $x_i \in X$ to a conjugate of itself. Theorem 5.4 of [7] establishes the isomorphism between the motion groups and the groups of automorphisms. If one does not restrict $H(S^3, L_{n+1})$ to those homeomorphisms preserving the orientations of L_{n+1} , then the resulting motion group is isomorphic to the subgroup of free group automorphisms sending each x_i to a conjugate of some $x_j^{\pm 1}$. The quotient of this larger motion group by the group Σ_{n+1} is the direct product of $(n+1)$ copies of \mathbb{Z}_2 .

The image of $P\Sigma_{n+1}$ in $\text{Aut}(F_{n+1})$ contains the inner automorphisms, hence one can form the quotient of $P\Sigma_{n+1}$ by the inner automorphisms, which we denote $\text{OP}\Sigma_{n+1}$. Our main cohomology argument is actually about $\text{OP}\Sigma_{n+1}$; Theorem 1.2 follows as a corollary.

In [13] McCullough and Miller introduced a family of contractible complexes that admit actions by certain automorphism groups of free products. They construct in particular a complex on which $\text{OP}\Sigma_n$ acts cocompactly. The fundamental domain for the action of $\text{OP}\Sigma_n$ on this space is isomorphic to the geometric realization of the poset HT_n . There is a bit of translation necessary since McCullough and Miller's presentation is in terms of bipartite trees and the Whitehead poset (see the remarks at the end of §2). We now define the complex and the action.

Definition 6.1 (Markings). A *marking* of a hypertree $\tau \in \text{HT}_{n+1}$ consists of a basis of F_{n+1} , which we denote $\{y_1, \dots, y_{n+1}\}$, where the element y_i is a conjugate of $x_i \in X$, and is associated with the vertex labelled i in τ .

Definition 6.2 (Marked automorphisms). An automorphism $\alpha \in P\Sigma_{n+1}$ is *carried* by a marked hypertree tree τ if:

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

Definition 6.3 (MM_{n+1}). The McCullough-Miller complex MM_{n+1} is the simplicial realization of the poset of marked hypertrees on $[n+1]$, modulo the equivalence relation generated by identifying two such trees if there is an automorphism carried by one that results in the other.

Theorem 6.4 (McCullough-Miller [13]). *The complex MM_{n+1} is a contractible complex of dimension $n-1$.*

The group $P\Sigma_{n+1}$ acts on MM_{n+1} by permuting the markings. The fundamental domain consists of the copy of $|\text{HT}_{n+1}|$ obtained by restricting to markings using the basis $X = \{x_1, \dots, x_{n+1}\}$. It is a *strong* fundamental domain in the sense that $P\Sigma_{n+1} \backslash MM_{n+1} \simeq |\text{HT}_{n+1}|$. Recalling the equivalence relation used in defining MM_{n+1} , it is clear that the group of inner automorphisms acts trivially on MM_{n+1} , hence the action of $P\Sigma_{n+1}$ yields an action of the quotient group $OP\Sigma_{n+1}$. In Section 5 of [13], McCullough and Miller compute the simplex stabilizers under the $OP\Sigma_{n+1}$ -action. The stabilizer of a vertex in MM_{n+1} , corresponding to a marked hypertree τ , is $\mathbb{Z}^{\text{rk}(\tau)}$, where $\text{rk}(\tau)$ is the number of hyperedges of τ minus one. More generally, the stabilizer of any simplex σ is the stabilizer of the vertex of smallest rank in σ . These results are recorded in the following lemma.

Lemma 6.5. *The fundamental domain for the action of $OP\Sigma_{n+1}$ on MM_{n+1} is isomorphic to the geometric realization $|\text{HT}_{n+1}|$. Viewing $\text{HT}_{n+1}^{[\infty]}$ as a subposet of HT_{n+1} one sees that the singular set of simplices whose isotropy group is a non-trivial finite rank free abelian group is $|\text{HT}_{n+1}^{[\infty]}|$. The stabilizer of any simplex in $|\text{HT}_{n+1}|$ that is not in $|\text{HT}_{n+1}^{[\infty]}|$ is trivial.*

7. BACKGROUND ON ℓ^2 -COHOMOLOGY

We only sketch the background necessary to follow our arguments; see [10] for additional background on ℓ^2 -cohomology.

Let G be the fundamental group of a compact aspherical complex X (or more generally an aspherical complex with finitely many cells in each dimension) so that the chain complex $C_*(\tilde{X})$ consists of finitely generated $\mathbb{Z}G$ -modules. If $\ell^2(G)$ is the Hilbert space of square-summable complex-valued functions on G , then $\ell^2(G)$ is a G -module and the induced ℓ^2 -cochain complex is

$$C^*(X, \ell^2(G)) = \text{Hom}_{\mathbb{Z}G}(C_*(\tilde{X}), \ell^2(G)) .$$

The orthogonal right action of G on $\ell^2(G)$ makes $C^*(X, \ell^2(G))$ a chain complex of Hilbert G -modules. If $\delta_i : C^i(X, \ell^2(G)) \rightarrow C^{i+1}(X, \ell^2(G))$ is the induced coboundary map then the *reduced ℓ^2 -cohomology groups of G* are defined to be


$$\mathcal{H}^i(G) = \text{Ker}(\delta_i) / \overline{\text{Im}(\delta_{i-1})}$$

where the overline indicates one takes the closure of the image. The classic example of a non-trivial ℓ^2 -cocycle is the cocycle for the trivalent tree indicated in Figure 5.

Each Hilbert G -module V has an associated von Neumann dimension $\dim_G(V)$ where, for example, $\dim_G(\underbrace{\ell^2(G) \oplus \dots \oplus \ell^2(G)}_{n \text{ copies}}) = n$. The ℓ^2 -Betti numbers of G

are then defined using the von Neumann dimension $b_*^{(2)}(G) = \dim_G \mathcal{H}^*(G)$. As was mentioned in the introduction, while the ℓ^2 -Betti numbers do not usually agree with ordinary Betti numbers, their alternating sum agrees with the ordinary Euler characteristic.

We use spectral sequences to calculate the ℓ^2 -Betti numbers of $P\Sigma_{n+1}$, hence we take a more algebraic perspective than what is described above. The *group von Neumann algebra* $\mathcal{N}(G)$ is the algebra of G -equivariant bounded operators from $\ell^2(G)$ to $\ell^2(G)$. As is discussed in §6 of Lück's book [10], the category of finitely generated Hilbert $\mathcal{N}(G)$ -modules is equivalent to the category of finitely generated projective $\mathcal{N}(G)$ -modules. For us, this has the benefit of allowing one to

FIGURE 5. An ℓ^2 -cocycle on the trivalent tree

use techniques from ordinary group cohomology (with coefficients in the group von Neumann algebra). In particular, we may appeal to the equivariant and Hochschild-Serre spectral sequences. In the end, one loses control of the functional analytic structure, but one can still compute the ℓ^2 -Betti numbers, since the definition of von Neumann dimension extends to arbitrary $\mathcal{N}(G)$ -modules (cf §6.5 in [10]). Further, as we discuss in the next section, this is sufficient for computing the ℓ^2 -cohomology of these groups.

One can define ℓ^2 -Betti numbers for arbitrary G -spaces, even spaces like MM_{n+1} where the action is not proper, by defining for any G -space X

$$b_p^{(2)}(X; \mathcal{N}(G)) = \dim_{\mathcal{N}(G)}(H_p^G(X; \mathcal{N}(G)))$$

where the singular homology $H_p^G(X; \mathcal{N}(G))$ is the homology of the $\mathcal{N}(G)$ -chain complex $\mathcal{N}(G) \otimes_{\mathbb{Z}G} C_*^{\text{sing}}(X)$.

The action of $\text{OP}\Sigma_{n+1}$ on MM_{n+1} has trivial or free abelian isotropy groups. Since the ℓ^2 -cohomology of \mathbb{Z}^k is trivial, the following result shows we may compute the ℓ^2 -Betti numbers of $\text{OP}\Sigma_{n+1}$ by computing the equivariant ℓ^2 -Betti numbers for this action. (The lemma is a consequence of [10, Theorem 6.54.2].)

Lemma 7.1. *Let X be a contractible G -CW-complex. Suppose that each isotropy group G_x is finite or satisfies $b_p^{(2)}(G_x) = 0$ for $p \geq 0$. Then $b_p^{(2)}(G) = b_p^{(2)}(X; \mathcal{N}(G))$ for $p \geq 0$.*

8. THE MAIN COHOMOLOGICAL ARGUMENT

Davis and Leary outline an approach to the following result in §10.1 of [5]; a more general result along these lines is proven by Davis, Januszkiewicz and Leary in §3 of [4].

Theorem 8.1. *Let G be a group of finite type admitting a cocompact action on a contractible complex X , with strong fundamental domain F . Let $X[\infty]$ be the subcomplex whose isotropy groups are infinite, and let $F[\infty] = F \cap X[\infty]$. Assume*

- (1) Each isotropy group G_x is trivial or satisfies $b_p^{(2)}(G_x) = 0$ for $p \geq 0$; and
- (2) The fundamental domain F is the cone over $F[\infty]$

Then the von Neumann dimensions of $H^i(X; \mathcal{N}(G))$ and $\mathcal{N}(G) \otimes \overline{H}^{i-1}(F[\infty])$ are the same. If in addition to (1) and (2),

- (3) G has no non-trivial element whose centralizer is finite index in G ,

then

$$\mathcal{H}^i(G) \simeq \ell^2(G) \otimes \overline{H}^{i-1}(F[\infty]) .$$

We briefly outline the Davis-Januszkiewicz-Leary approach to this result. Given conditions (1) and (2) on the action of G on X one can form a complex of spaces Y based on the fundamental domain F . One can then apply a Leray, or generalized Mayer-Vietoris, spectral sequence to compute cohomology. A careful analysis of the resulting spectral sequence applied when the module is $\mathcal{N}(G)$ establishes that the von Neumann dimensions of $H^i(X, \mathcal{N}(G))$ and $\overline{H}^{i-1}(F[\infty]) \otimes \mathcal{N}(G)$ are the same. The dimension shift that occurs in this formula is natural, as one is essentially working with the cohomology of the pair $H^i(|\mathrm{HT}_{n+1}|, |\mathrm{HT}_{n+1}^{[\infty]}|)$; by the long exact sequence in cohomology this is isomorphic to $H^{i-1}(|\mathrm{HT}_{n+1}^{[\infty]}|)$ since $|\mathrm{HT}_{n+1}|$ is contractible. If condition (3) holds, then the von Neumann dimension of a finitely generated Hilbert G -module determines the module (see §2.1 and 2.2 of [14] and §9.1 in [10]), hence the concluding isomorphism.

Theorem 8.2. *The ℓ^2 -cohomology of $\mathrm{OP}\Sigma_{n+1}$ is described by*

$$\mathcal{H}^i(\mathrm{OP}\Sigma_{n+1}) \simeq \ell^2(\mathrm{OP}\Sigma_{n+1}) \otimes \overline{H}^{i-1}(|\mathrm{HT}_{n+1}^{[\infty]}|) .$$

In particular, the ℓ^2 -Betti numbers of $\mathrm{OP}\Sigma_{n+1}$ are all trivial except in top dimension, where

$$\chi(\mathrm{OP}\Sigma_{n+1}) = (-1)^{n-1} b_{n-1}^{(2)} = (-1)^{n-1} n^{n-1} .$$

Proof. It follows from the description in Lemma 6.5 that action of $\mathrm{OP}\Sigma_{n+1}$ on MM_{n+1} satisfies conditions (1) and (2) of Theorem 8.1. Thus we may compute the ℓ^2 -Betti numbers of $\mathrm{OP}\Sigma_{n+1}$. Since $|\mathrm{HT}_{n+1}^{[\infty]}|$ is Cohen-Macaulay (Theorem 2.9), its cohomology is concentrated in top dimension. It follows that the ordinary Betti numbers of $|\mathrm{HT}_{n+1}^{[\infty]}|$ are trivial, except in top dimension, where the Betti number is the absolute value of the reduced Euler characteristic $\tilde{\chi}(|\mathrm{HT}_{n+1}^{[\infty]}|)$. Thus the ℓ^2 -Betti number $b_{n-1}^{(2)}(\mathrm{OP}\Sigma_{n+1})$ equals the ordinary Betti number $b_{n-1}(|\mathrm{HT}_{n+1}^{[\infty]}|)$, and the reduced Euler characteristic formula of Theorem 1.1 implies the formula above.

Given the material presented so far one can best demonstrate that condition (3) of Theorem 8.1 holds by noting first that $\mathrm{OP}\Sigma_{n+1}$ is torsion free. If the centralizer of some $g \in \mathrm{OP}\Sigma_{n+1}$ ($g \neq 1$) was of finite index in $\mathrm{OP}\Sigma_{n+1}$, then the ℓ^2 -Betti numbers of $\mathrm{OP}\Sigma_{n+1}$ would be trivial. This follows since the ℓ^2 -Betti numbers of $C(g)$ would be trivial (Corollary 8 in [5]) and ℓ^2 -Betti numbers are multiplicative. But we already know that $b_{n-1}^{(2)}(\mathrm{OP}\Sigma_{n+1}) = n^{n-1}$, hence (3) follows. \square

If one defines $\mathrm{O}\Sigma_{n+1}$ to be the quotient of Σ_{n+1} by the inner-automorphisms, it follows immediately from the fact that ℓ^2 -Betti numbers are multiplicative that

the ℓ^2 -Betti numbers of $O\Sigma_{n+1}$ are also trivial except in top dimension, where they are described by the amusing formula

$$b_{n-1}^{(2)}(O\Sigma_{n+1}) = \frac{n^{n-1}}{(n+1)!}.$$

We can now prove Theorem 1.2 using the short exact sequence

$$1 \rightarrow F_{n+1} \rightarrow P\Sigma_{n+1} \rightarrow OP\Sigma_{n+1} \rightarrow 1.$$

Since the ℓ^2 -cohomology of the kernel and quotient are concentrated in top dimension, it follows from the Hochschild-Serre spectral sequence that the ℓ^2 -cohomology of $P\Sigma_{n+1}$ is concentrated in top dimension. (This use of the Hochschild-Serre spectral sequence was inspired by the proof of 6.66 and 6.67 in [10].) Since the Euler characteristic of $P\Sigma_{n+1}$ is the product of the Euler characteristics of F_{n+1} and $OP\Sigma_{n+1}$ [2, Prop 7.3], and $\chi(F_{n+1}) = -n$, it follows that $\chi(P\Sigma_{n+1}) = (-1)^n n^n$. But because the alternating sum of the ℓ^2 -Betti numbers gives the ordinary Euler characteristic, it follows that

$$\chi(P\Sigma_{n+1}) = (-1)^n b_n^{(2)}(P\Sigma_{n+1}) = (-1)^n n^n.$$

REFERENCES

- [1] 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.
- [2] Kenneth S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.
- [3] Donald J. Collins. Cohomological dimension and symmetric automorphisms of a free group. *Comment. Math. Helv.*, 64(1):44–61, 1989.
- [4] M. W. Davis, T. Januszkiewicz, and I. J. Leary. The ℓ^2 -cohomology of hyperplane complements. in progress.
- [5] M. W. Davis and I. J. Leary. The ℓ^2 -cohomology of Artin groups. *J. London Math. Soc. (2)*, 68(2):493–510, 2003.
- [6] Jan Dymara and Tadeusz Januszkiewicz. Cohomology of buildings and their automorphism groups. *Invent. Math.*, 150(3):579–627, 2002.
- [7] Deborah L. Goldsmith. The theory of motion groups. *Michigan Math. J.*, 28(1):3–17, 1981.
- [8] Mauricio Gutiérrez and Sava Krstić. Normal forms for basis-conjugating automorphisms of a free group. *Internat. J. Algebra Comput.*, 8(6):631–669, 1998.
- [9] L. Kalikow. *Enumeration of parking functions, allowable permutation pairs, and labeled trees*. PhD thesis, Brandeis University, 1999. Available at <http://home.gwu.edu/~lkalikow/research/research.html>.
- [10] Wolfgang Lück. *L^2 -invariants: theory and applications to geometry and K-theory*, volume 44 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics*. Springer-Verlag, Berlin, 2002.
- [11] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Classics in Mathematics. Springer-Verlag, Berlin, 2001. Reprint of the 1977 edition.
- [12] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial group theory*. Dover Publications Inc., New York, revised edition, 1976.
- [13] Darryl McCullough and Andy Miller. Symmetric automorphisms of free products. *Mem. Amer. Math. Soc.*, 122(582):viii+97, 1996.
- [14] Shōichirō Sakai. *C^* -algebras and W^* -algebras*. Classics in Mathematics. Springer-Verlag, Berlin, 1998. Reprint of the 1971 edition.
- [15] Richard P. Stanley. *Enumerative combinatorics. Vol. 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997.
- [16] Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999.

- [17] D. Warne. *Spanning trees in hypergraphs with applications to Steiner trees*. PhD thesis, University of Virginia, 1998. Available at <http://s3i.com/~warne/pubs/>.

DEPT. OF MATH., U. C. SANTA BARBARA, 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