COMBINATORIAL LAPLACIANS OF MATROID COMPLEXES

W. KOOK, V. REINER, AND D. STANTON

Abstract. We combinatorially interpret the spectra of discrete Laplace operators from the boundary maps in the simplicial complex of independent sets of a matroid. The interpretation follows from a surprising orthogonal decomposition of the simplicial chain groups. This decomposition is in general finer than the spectral decomposition. As a consequence, the spectra are integral.

One corollary to our combinatorial interpretation may be paraphrased as stating that one can "hear" the characteristic polynomial of a matroid.

1. Introduction

For any finite simplicial complex K, one can define Laplace operators Δ_i which are combinatorial analogues of the Laplace operators on differential forms for a Riemannian manifold. The definition (as in [5]) is as follows. Let C_i be the $\mathbb{R}$-vector space of (oriented) simplicial i-chains in K with real coefficients, and $\partial_i : C_i \to C_{i-1}$ the usual simplicial boundary map. Endow C_i with an inner product by declaring the oriented simplices to form an orthonormal basis of C_i, so that we may identify C_i with its dual C_i^*. The adjoint to ∂_i with respect to this inner product is the transpose ∂_i^T or the coboundary map $\delta_{i-1} : C_{i-1} \to C_i$. Define $\Delta_i : C_i \to C_i$ by

$$\Delta_i = \delta_{i-1}\partial_i + \partial_{i+1}\delta_i.$$

The main result about combinatorial Laplacians is an analogue of a fact from Hodge theory, namely that the kernel or 0-eigenspace of Δ_i is naturally isomorphic to the homology group $H_i(K;\mathbb{R})$ [5, Proposition 2.1]. The computational advantage of considering Δ_i is that one can compute numerical approximations to the spectrum of Δ_i and find its 0-eigenspace relatively quickly, using well-established fast methods for spectra of real symmetric matrices [5].

These methods are particularly effective if one knows that the spectra of the Δ_i are integral, which is not often the case. However, the spectra are integral for the well-studied *chessboard complexes*, as shown in [6]. There the authors used the symmetry groups acting on the complexes for a clever representation-theoretic calculation, which gives the spectrum and in particular its integrality.

The main result of this paper (Corollary 17) proves that the spectra of all the Laplacians are again integral when K is a *matroid complex*, that is when K is the complex of independent sets in a matroid. The result follows from a surprising orthogonal decomposition of the chain groups C_i (Theorem 15) which is in general finer than the eigenspace decomposition for the Laplacians. We remark that even though the combinatorial Laplacians have usually been defined only over fields of characteristic zero, all of our results are still true over the integers.

Key words and phrases. matroid, matroid complex, Laplacian, internal activity, external activity.

Second author supported by Sloan Foundation and Univ. of Minnesota McKnight Land Grant Fellowships. Third author supported by NSF grant DMS-9400510.

One corollary (Corollary 10) to our results is that for a matroid complex, only a very small part of the information in the spectra of the Laplacians (namely the multiplicity of the largest eigenvalue in each Δ_i) is sufficient to recover the *characteristic polynomial* of the matroid. This adds to the already large list (see e.g. [3]) of interpretations of the characteristic polynomial of a matroid, and can be whimsically rephrased by saying that one can "hear" the characteristic polynomial of a matroid.

The paper is structured as follows. Section 2 deals with some combinatorial preliminaries that do not involve Laplacians, but are really the combinatorial manifestations that foreshadow Theorem 15. Section 3 introduces the combinatorial Laplacians and proves some easy basic facts about their behavior for matroid complexes. Section 4 contains the main results, which require a slightly deeper analysis, aided by the technique of *exterior face rings* for simplicial complexes [4]. Section 5 discusses some well-known examples of matroids having large symmetry groups.

2. Combinatorial preliminaries

We refer the reader to the excellent article [1] for almost all of the definitions and notions from matroid theory and matroid complexes that we will use.

Let M be a loopless matroid with a chosen linear ordering ω on its ground set E. In the usual way there is associated to M a *rank function* $r(A)$ and a *closure operator* $A \mapsto \overline{A}$ on subsets $A \subseteq E$, and the set of closed subsets or *flats* forms a *geometric lattice* denoted $L(M)$.

Recall that for a base B of M, an element b in B of M gives rise to a unique *basic bond* $bo_M(B,b)$, and an element e in $E-B$ gives rise to a unique *basic circuit* $ci_M(B,e)$. We say that b is *internally active* in B if it is the ω-smallest element of $bo_M(B,b)$, and e is *externally active* if it is the ω-smallest element of $ci_M(B,e)$. The *internal activity* $i(B)$ of the base B is the number of internally active elements in B, and its *external activity* $e(B)$ is defined similarly.

Theorem 1. *In any ordered matroid (M,ω), every base B has a unique disjoint decomposition $B = B_1 \sqcup B_2$ with the following properties:*

- *B_1 is a base of internal activity 0 for the flat $V := \overline{B_1}$ which it spans, and*
- *B_2 is a base of external activity 0 for the quotient matroid M/V.*

Proof. The existence of such a decomposition is provided by the following algorithm, which takes the base B as input, and produces such a pair (B_1, B_2).

ALGORITHM:

Step 1. Set $B_1 = B, B_2 = \varnothing$.

Step 2. Let $V = \overline{B_1}$.

Step 3. Find an internally active element b for B_1 as a base of the flat V.

If no such element b exists, then stop and output the pair (B_1, B_2).

If such a b exists, then set $B_1 := B_1 - \{b\}, B_2 := B_2 \cup \{b\}$, and return to Step 2.

It is clear that the algorithm terminates, since each time one returns to Step 2, the cardinality of B_1 has decreased. To see that the output is a decomposition as in the theorem, note that B_1 will be a base for V of internal activity 0 by construction, so the first property is clearly satisfied.

Therefore it remains to show the second property, i.e. that B_2 is a base of external activity 0 for M/V. For this we will show that given any $e \in E - V - B_2$

the circuit $ci_{M/V}(B_2, e)$ contains an ω-smaller element than e. Since $ci_{M/V}(B_2, e)$ is a circuit in M/V, the set $ci_{M/V}(B_2, e) \cup B_1$ is dependent in M, and therefore contains some circuit of M of the form $ci_{M/V}(B_2, e) \cup B_1'$ where $B_1' \subseteq B_1$. Denote by $(B_1^{(i)}, B_2^{(i)})$ the sets (B_1, B_2) at the i^{th} stage of the above algorithm, so $B_1^{(0)} = B, B_2^{(0)} = \varnothing$. Let i be such that $ci_{M/V}(B_2, e) - \{e\} \subseteq B_1^{(i)}$ but $ci_{M/V}(B_2, e) - \{e\} \not\subseteq B_1^{(i+1)}$. This means that some element b_2 of $ci_{M/V}(B_2, e) - \{e\}$ is thrown out of $B_1^{(i)}$ and into $B_2^{(i+1)}$ at this stage because b_2 is internally active for $B_1^{(i)}$. Note that $ci_{M/V}(B_2, e) - \{e\} \cup B_1' \subseteq B_1^{(i)}$ and consequently $b_2 \in ci_M(B^{(i)}, e)$ with b_2 being ω-smaller than e.

To show uniqueness of the decomposition, we define a map f between two sets:

- the set of triples (V, B_1, B_2) where V is a flat of M, B_1 is a base of internal activity 0 for V, and B_2 is a base of external activity 0 for M/V, and
- the set of bases B of M.

The map f sends the triple (V, B_1, B_2) to the base $B := B_1 \sqcup B_2$. Uniqueness of the decomposition is equivalent to injectivity of f, and since the above algorithm shows that f is surjective, it suffices to show that the two sets have the same cardinality, i.e. that

(2.1)

$$| \text{ bases of M } | = \sum_{\text{flats } V \text{ of } M} \left| \begin{array}{c} \text{bases } B_1 \text{ of } V \\ \text{with internal activity 0} \end{array} \right| \times \left| \begin{array}{c} \text{bases } B_2 \text{ of } M/V \\ \text{with external activity 0} \end{array} \right|.$$

To see this, recall from [1, Theorem 7.8.4] that the number of bases of a matroid M having internal activity 0 is, up to sign, the (reduced) Euler characteristic $\tilde{\chi}(IN(M))$ of the *independence complex* $IN(M)$, and hence is equal to the alternating sum $\sum_I (-1)^{r(M)-|I|}$ where I ranges over the independent sets of M. Recall also from [1, Proposition 7.4.7] that the number of bases of a matroid M having external activity 0 is the absolute value of the Möbius function $|\mu_{L(M)}(\hat{0}, \hat{1})|$ for the geometric lattice $L(M)$. Therefore starting with the right-hand side in equation (2.1), we can manipulate as follows:

$$\begin{aligned}
&\sum_{\text{flats } V \text{ of } M} \left| \begin{array}{c} \text{bases } B_1 \text{ of } V \\ \text{with internal activity 0} \end{array} \right| \times \left| \begin{array}{c} \text{bases } B_2 \text{ of } M/V \\ \text{with external activity 0} \end{array} \right| \\
&= \sum_{\text{flats } V \text{ of } M} \left(\sum_{\text{independent sets } I \subseteq V} (-1)^{r(V)-|I|} \right) |\mu_{L(M/V)}(\hat{0}, \hat{1})| \\
&= \sum_{\text{independent sets } I \subseteq E} (-1)^{r(M)-|I|} \sum_{\overline{I} \leq V \leq \hat{1}} (-1)^{r(M)-r(V)} |\mu_{L(M)}(V, \hat{1})| \\
&= \sum_{\text{independent sets } I \subseteq E} (-1)^{r(M)-|I|} \sum_{\overline{I} \leq V \leq \hat{1}} \mu_{L(M)}(V, \hat{1}) \\
&= \sum_{\text{independent sets } I \text{ with } \overline{I} = M} (-1)^{r(M)-|I|} \\
&= \sum_{\text{bases } B \text{ of } M} (-1)^{r(M)-|B|} \\
&= | \text{ bases of M } |.
\end{aligned}$$

Here we have used various facts, such as the defining recurrence for the Möbius function, the fact that the Möbius function of the geometric lattice $L(M)$ alternates in sign, and the fact that the interval $[V, \hat{1}|$ in $L(M)$ is isomorphic to the lattice of flats $L(M/V)$ for the quotient matroid M/V. □

Based on Theorem 1, we define the *spectrum polynomial* of M

$$\mathrm{Spec}_M(t,q) := \sum_{\text{independent sets } I} t^{r(\overline{I})} q^{|I_1|}$$

where I_1 denotes the output of the decomposition algorithm $I \mapsto (I_1, I_2)$ when I is considered as a base for the flat $\overline{I}$ it spans, and the ordering on the ground set in $\overline{I}$ is induced from the original ordering ω on E.

The reason for the terminology "spectrum polynomial" is that $\mathrm{Spec}_M(t,q)$ carries the same information as the spectra of the Laplacians for the matroid complex of M (Corollary 17). The simpler fact that $\mathrm{Spec}_M(t,q)$ does not depend on the chosen ordering ω of the ground set is easy to see from the following alternate expression:

$$\begin{aligned} \mathrm{Spec}_M(t,q) &= \sum_{\text{flats } V} t^{r(V)} \sum_{\text{independent sets } I \text{ spanning } V} q^{|I_1|} \\ &= \sum_{\text{flats } V} t^{r(V)} \sum_{\text{flats } V' \subseteq V} |\tilde{\chi}(IN(V'))| \cdot |\mu_{L(M)}(V', V)| q^{|V'|} \end{aligned} \tag{2.2}$$

The last equality follows from the same facts used to show that the map f was a bijection in the proof of Theorem 1.

A straightforward consequence of the definition is that for a direct sum $M_1 \oplus M_2$ of two matroids M_1, M_2, we have

$$\mathrm{Spec}_{M_1 \oplus M_2}(t,q) = \mathrm{Spec}_{M_1}(t,q) \cdot \mathrm{Spec}_{M_2}(t,q), \tag{2.3}$$

a property reminiscent of the Tutte polynomial of a matroid. Recall that the *Tutte polynomial* $T_M(x,y)$ may be defined by

$$T_M(x,y) := \sum_{\text{bases } B \text{ of } M} x^{i(B)} y^{e(B)}$$

where $i(B), e(B)$ are the internal and external activity of the base B. The Tutte polynomial plays an important role in matroid theory; see [3]. Since the Tutte polynomial and spectrum polynomial are both related to the notions of internal and external activity, it is tempting to compare them.

Question 2. *How does the spectrum polynomial compare with the Tutte polynomial in distinguishing non-isomorphic matroids? Can two matroids have the same spectrum polynomial and different Tutte polynomials?*

We know that there are non-isomorphic matroids that have same spectrum polynomial. For example, the matroids M_1, M_2 represented in characteristic zero by the affine point configurations shown in Figure 1 have the same spectrum polynomial

$$1 + 6t + t^2(13 + 2q^4) + t^3(8 + 2tq^3 + 8q^6)$$

and the same Tutte polynomial [3, Problem 6.1(b)]

$$x^3 + y^3 + 3x^2 + 3y^2 + 3xy + 4x + 4y.$$

On the other hand, there are non-isomorphic matroids that have the same Tutte

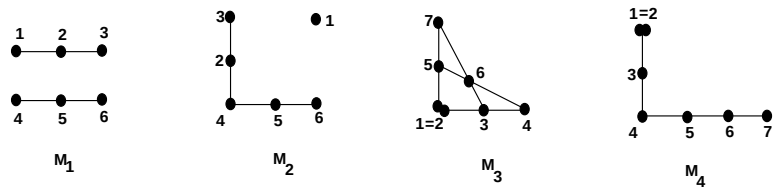

FIGURE 1. Matroids M_1 and M_2 have the same Tutte polynomial and the same spectrum polynomial. Matroids M_3 and M_4 have the same Tutte polynomial but different spectrum polynomials.

polynomial, but not the same spectrum polynomial. For example, the matroids M_3, M_4 shown in Figure 1 have the same Tutte polynomial [3, Example 6.2.18]

$$y^4 + 3y^3 + 2xy^2 + x^2y + x^3 + 4y^2 + 5xy + 3x^2 + 2x + 2y,$$

but different spectrum polynomials

$$\begin{aligned}\mathrm{Spec}_{M_3}(t,q) =& 1 + t(6+q^2) + t^2(11+3q^2+2q^3+4q^4) + t^3(6+2q^2+2q^3+4q^4+10q^7)\\ \mathrm{Spec}_{M_4}(t,q) =& 1 + t(6+q^2) + t^2(11+4q^2+5q^4) + t^3(6+3q^2+5q^4+10q^7).\end{aligned}$$

We see no reason to expect that the spectrum polynomial determines the Tutte polynomial, but we have no counterexamples.

Since the Tutte polynomial is easy to compute by the *deletion-contraction recurrence*

$$T_M(x,y) = T_{M-e}(x,y) + T_{M/e}(x,y),$$

the following question is natural:

Question 3. *Is there some way to use deletion-contraction to compute the spectrum polynomial* $\mathrm{Spec}_M(t,q)$*?*

The examples M_3, M_4 above show that there can be no formula for $\mathrm{Spec}_M(t,q)$ purely as a function of $\mathrm{Spec}_{M-e}(t,q), \mathrm{Spec}_{M/e}(t,q)$, since $M_3 - e \cong M_4 - e$ and $M_3/e \cong M_4/e$ for the element e represented by the 7^{th} column of the matrices in each case. Nevertheless, there might exist some more complicated procedure for computing the spectrum polynomial of M from those of a deletion $M-e$ and contraction M/e, which somehow uses the way in which $M-e$ and M/e "glue together" to produce M.

3. Laplacians

The goal of this section is to define the combinatorial Laplacians of matroid complexes, and prove some easy facts about them.

Let $IN(M)$ denote the independence complex of the matroid M, i.e. the simplicial complex on vertex set E whose simplices are the independent sets of M. Let C_i^M denote the i^{th} (oriented) simplicial chain group of $IN(M)$ with $\mathbb{R}$ coefficients, $\partial_i^M : C_i^M \to C_{i-1}^M$ the usual simplicial boundary map, and $\tilde{H}_i^M$ the i^{th} (reduced) homology group with $\mathbb{R}$ coefficients. By choosing the standard inner product $\langle \cdot, \cdot \rangle$ on C_i^M in which the oriented simplices are an orthonormal basis, we can identify C_i^M with its dual space of cochains. Then the simplicial coboundary map δ_i^M

which is the adjoint of ∂_i^M may be thought of as a map $C_{i-1}^M \to C_i^M$. Define the i^{th} combinatorial Laplacian to be the operator $\Delta_i^M : C_i^M \to C_i^M$ given by

$$\Delta_i^M := \delta_{i-1}^M \partial_i^M + \partial_{i+1}^M \delta_i^M.$$

It is not hard to show [5] that the kernel or 0-eigenspace $ker(\Delta_i^M)$ is isomorphic to the homology $\tilde{H}_i^M$. In our case, there is not just an isomorphism but an *equality* $ker(\Delta_{r(M)-1}^M) = \tilde{H}_{r(M)-1}^M$ since there are no $(r(M)-1)$-dimensional boundaries. Furthermore, since $IN(M)$ is shellable [1, Theorem 7.3.3] and hence has only $(r(M)-1)$-dimensional homology, we have that $ker(\Delta_i^M) = 0$ for $i < r(M)-1$, a fact we will recover for different reasons shortly (Corollary 7).

It is easy to compute the following explicit description for Δ_i^M. For each independent $(i+1)$-set I of M, if we let $[I]$ denote the oriented simplex which places the elements of I in ω-order, then

$$\Delta_i^M [I] = (|I| + |E - \overline{I}|) \cdot [I] + \sum_{\substack{I' = I - \{e\} + \{e'\} \\ I' \cup I \text{ dependent}}} (-1)^{d(I,I')} [I'] \tag{3.1}$$

where $d(I, I')$ is the number of elements in $I \cup I'$ which lie strictly between e, e' under the ω-ordering.

Remark 4. Apparently Δ_i^M depends upon on ω. However, one can check that changing ω only conjugates Δ_i^M by a signed permutation matrix.

First a few simple observations.

Proposition 5. *The chain group C_i^M and the operator Δ_i^M decompose as a direct sum*

$$C_i^M = \bigoplus_{rank\ (i+1)\ flats\ V} C_i^V$$

$$\Delta_i^M = \bigoplus_{rank\ (i+1)\ flats\ V} \left(\Delta_i^V + |E - V| \cdot I_V\right).$$

Here C_i^V is the subspace of i-chains supported on independent sets I which are a base for the flat V, and I_V is the identity map on this subspace, while Δ_i^V is the i^{th} Laplacian for the restricted matroid $M|_V$.

Proof. Looking at the explicit formula (3.1) for Δ_i^M, one sees that the (I, I') off-diagonal entry is non-zero only when I and I' are bases of the same flat V. Furthermore, in this case the scalars $|E - \overline{I}|, |E - \overline{I'}|$ are both equal to $|E - V|$. □

As a consequence of the previous proposition, we really need only understand the top Laplacian $\Delta_{r(M)-1}^M$ acting on $C_{r(M)-1}^M$.

To this end, recall that the map $B \mapsto E - B$ gives a bijection between the bases for M and the bases for its dual matroid M^*. Extend this bijection to a linear isomorphism

$$P : C_{r(M)-1}^M \to C_{r(M^*)-1}^{M^*}$$

by sending the oriented simplex $[B]$ to $(-1)^{\text{ranksum}(B)}[E - B]$ where ranksum(B) is the sum of the ranks of the elements of B in the ω-order.

Proposition 6.

$$\Delta^M_{r(M)-1} \quad + \quad P^{-1}\Delta^{M^*}_{r(M^*)-1}P \quad = \quad |E| \cdot I_M$$

where I_M is the identity map on $C^M_{r(M)-1}$.

Proof. Examining the explicit formula (3.1) for Δ^M_i in the case $i = r(M)-1$, we see that the diagonal entries for $\Delta^M_{r(M)-1}, P^{-1}\Delta^{M^*}_{r(M^*)-1}P$ are $r(M), |E| - r(M)$ respectively. Consequently the assertion for the diagonal entries in the proposition is clear. For the off-diagonal entries, one needs to check that for any bases B, B' of M satisfying $B' = B - \{e\} \cup \{e'\}$ we have

$$(-1)^{\mathrm{ranksum}(B)}(-1)^{\mathrm{ranksum}(B')}(-1)^{d(B,B')} + (-1)^{d(E-B,E-B')} = 0.$$

This verification is straightforward. □

Corollary 7. *For every i, the spectrum of Laplacian Δ_i is contained in the interval $[0, |E|]$ of $\mathbb{R}$, and will not contain 0 unless $i = r(M)-1$.*

Proof. We know that Δ^M_i is non-negative definite since it is a sum of the two operators $\delta^M_{i-1}\partial^M_i$ and $\partial^M_{i+1}\delta^M_i$, which are both non-negative definite since they are of the form A^*A (see [5, Proposition 2.1]). Therefore the spectrum of Δ^M_i is nonnegative.

If we look at the special case of $i = r(M)-1$, then the previous proposition says that the spectrum of $\Delta^M_{r(M)-1}$ is obtained from the (nonnegative) spectrum of $\Delta^{M^*}_{r(M^*)-1}$ by subtraction from $|E|$. Therefore the assertion of the corollary holds for $i = r(M)-1$. But then a glance at Proposition 5 shows that it also holds for the remaining i. □

Proposition 6 shows that the spectrum of the top Laplacian $\Delta^M_{r(M)-1}$ completely determines the spectrum of the dual $\Delta^{M^*}_{r(M^*)-1}$. This suggests the following question:

Question 8. *Is the same true for the lower-dimensional Laplacians Δ^M_i, or do there exist two matroids M_1, M_2 for which $\Delta^{M_1}_i$ and $\Delta^{M_2}_i$ have the same spectra for all i, but some $\Delta^{M_1^*}_j$ and $\Delta^{M_2^*}_j$ have different spectra?*

This also suggests a stronger question to which we also do not know the answer:

Question 9. *To what extent does the top spectrum determine the lower spectra?*

For a real number λ, let $(\Delta^M_i)_\lambda$ denote the λ-eigenspace for the operator Δ^M_i. Corollary 7 shows that the largest possible value of λ is the cardinality of the ground set $|E|$. Our next result gives an interesting interpretation for the multiplicities of this largest eigenvalue in the various Laplacians Δ^M_i.

Recall that the characteristic polynomial of the matroid is defined by

$$\chi_M(t) := \sum_{V \in L(M)} \mu_{L(M)}(\hat{0}, V)\, t^{r(M)-r(V)}$$

where $\mu_{L(M)}$ is the Möbius function for the geometric lattice $L(M)$. The next corollary has the whimsical interpretation that you can "hear" the characteristic polynomial of M, if by analogy with vibrating drums, we imagine that the spectra of the various Laplacians Δ^M_i are heard in the frequencies emitted when one "bangs" a matroid complex!

Corollary 10.

$$\chi_M(-t) = t^{r(M)-1} \sum_{i=-1}^{r(M)-1} \dim_{\mathbb{R}}(\Delta_i^M)_{|E|} t^{-i}$$

Proof.

$$\begin{aligned}
&(-t)^{r(M)-1} \sum_{i=-1}^{r(M)-1} \dim_{\mathbb{R}}(\Delta_i^M)_{|E|}(-t)^{-i} \\
&= (-t)^{r(M)-1} \sum_{i=-1}^{r(M)-1} \sum_{\text{flats } V: r(V)=i+1} \dim_{\mathbb{R}}(\Delta_i^V)_{|E|}(-t)^{-i} \\
&= \sum_{\text{flats } V} \dim_{\mathbb{R}}(\Delta^V_{r(V)-1})_{|E|}\,(-t)^{r(M)-r(V)} \\
&= \sum_{\text{flats } V} \dim_{\mathbb{R}}(\Delta^{V^*}_{r(V^*)-1})_0\,(-t)^{r(M)-r(V)} \\
&= \sum_{\text{flats } V} \dim_{\mathbb{R}} \tilde{H}_{r(V^*)-1}(IN(V^*))(-t)^{r(M)-r(V)} \\
&= \sum_{\text{flats } V} |\mu_{L(V)}(\hat{0},\hat{1})|(-t)^{r(M)-r(V)} \\
&= \sum_{V \in L(M)} \mu_{L(M)}(\hat{0}, V) t^{r(M)-r(V)}
\end{aligned}$$

The first equality uses Proposition 5. The second equality is just interchanging the order of summation. The third equality uses Proposition 6. The fourth equality uses the fact the 0-eigenspace for the top Laplacian is the same as the top homology of the independence complex. The fifth equality uses the fact from [1, Theorem 7.8.1] that this top homology has dimension equal to Möbius number for the geometric lattice of the dual matroid. The sixth equality uses the fact that the Möbius function for the geometric lattice $L(M)$ alternates in sign. □

It is well-known that the Tutte-polynomial $T_M(x, y)$ also specializes to the characteristic polynomial:

$$\chi_M(t) = (-1)^{r(M)} T_M(1 - t, 0).$$

The generating function

$$\sum_{\text{independent sets } I \text{ in } M} t^{|I|} = t^{r(M)}\, T_M(1 + t^{-1}, 1)$$

is another specialization of the Tutte polynomial which is clearly recoverable from the spectra of the various Δ_i^M. This prompts the following question:

Question 11. *How much can one "hear" of the Tutte polynomial of a matroid, i.e. how much of $T_M(x, y)$ is recoverable from knowledge of the spectra of the various Laplacians Δ_i^M ?*

This question will turn out to be equivalent to Question 2 via Corollary 17.

4. Orthogonal decomposition of $C^M_{r(M)-1}$ and the exterior face ring

The goal of this section is to probe deeper into the spectra of the Laplacians Δ^M_i, which by Proposition 5 reduces to understanding $\Delta^M_{r(M)-1}$. We will show that all of its eigenvalues are integers, and explain how the eigenspace decomposition of $C^M_{r(M)-1}$ with respect to $\Delta^M_{r(M)-1}$ is refined by a more fundamental orthogonal decomposition (Theorem 15) whose combinatorial manifestation is Theorem 1. As a consequence, we will deduce our main result (Corollary 17) that the spectrum polynomial $\mathrm{Spec}_M(t,q)$ is essentially the generating function for the spectra of $\Delta^M_i, i = -1, \ldots, r(M)-1$.

Our primary tool will be the *exterior face ring* $\wedge M$ of $IN(M)$ [4]. Let W be an $\mathbb{R}$-vector space with basis $\{v_e\}_{e\in E}$ indexed by the ground set E of M. Let $\wedge W$ be the exterior algebra on W, having basis $v_S := v_{e_1} \wedge \cdots \wedge v_{e_s}$ indexed by subsets $S = \{e_1, \ldots, e_s\} \subseteq E$. The exterior face ring $\wedge M$ is the quotient

$$\wedge M := \wedge W/(v_S : S \text{ is dependent in } M).$$

As an $\mathbb{R}$-vector space, we will identify $\wedge M = \bigoplus_{i=0}^{r(M)} \wedge^i M$ with the direct sum of the oriented chain groups $\bigoplus_{i=-1}^{r(M)-1} C^M_i$. As a consequence, this direct sum of chain groups is endowed with an $\mathbb{R}$-algebra structure; under this identification the coboundary maps δ_i correspond to multiplication on the left by the fixed element of degree one $\delta^M := \sum_{e\in E} v_e$. Furthermore, for any subset $S \subseteq E$, let δ^S denote the element of degree one $\sum_{s\in S} v_s$. The inner product on $\bigoplus_i C^M_i$ corresponds to one on $\wedge M$ in which the $\{v_I\}$ are declared to form an orthonormal basis. With respect to this inner product, the boundary maps ∂_i may be thought of as a single map ∂^M which is adjoint to left multiplication by δ^M, i.e.

$$\langle x, \partial^M y\rangle = \langle \delta^M \wedge x, y\rangle.$$

One crucial (and easy to check) property of ∂^M is that it is almost a (graded) derivation on the algebra $\wedge M$. We have

$$\partial^M(x \wedge y) = \partial^M(x) \wedge y + (-1)^{|x|} x \wedge \partial^M(y) \tag{4.1}$$

for any $x \in \wedge^{|x|} M$, and y having the property that *no* v_S, v_T appearing in x, y respectively have $v_S \wedge v_T = 0$.

We will mostly be dealing with the top dimensional Laplacian, which for shorter notation we will denote Δ^M rather than $\Delta^M_{r(M)-1}$. Similarly let C^M replace the notation $C^M_{r(M)-1}$ for the top dimensional chain group. Notice that since there is no chain group in dimension $r(M)$ we have

$$\Delta^M(x) = \delta^M \wedge \partial^M(x)$$

for all $x \in \wedge^{r(M)} M$. As a consequence of this fact and our use of $\mathbb{R}$ coefficients, $H^{r(M)-1}(IN(M))$ is a subspace of the 0-eigenspace $(\Delta^M)_0$. In fact these two spaces are equal since they have the same dimension.

Note that if V is a flat of M then we can consider V as a matroid by restriction, and $\wedge V$ is naturally a subalgebra of $\wedge M$, so we can think of elements of $\wedge V$ as being in $\wedge M$. Also $\wedge M/V$ is naturally a quotient of $\wedge M$, and by abuse of notation, we can consider elements of $\wedge M/V$ as elements of $\wedge M$ in the following way: if $S \subseteq E - V$ is independent in M/V then identify $v_S \in \wedge M/V$ with the element $v_S \in \wedge M$.

Lemma 12. *If $x \in \wedge^{r(V)}V$ and $y \in \wedge^{r(M/V)}M/V$ satisfy*

$$\Delta^V(x) = \lambda \cdot x$$
$$\Delta^{M/V}(y) = \mu \cdot y$$

for some real constants λ, μ, then

$$\Delta^M(x \wedge y) = (\lambda + \mu) \cdot x \wedge y + \delta^{E-V} \wedge \partial^M(x) \wedge y$$

Proof. Starting with the left-hand side, we perform a series of manipulations which are justified below

$$\begin{aligned}
&\Delta^M(x \wedge y) \\
&= \delta^M \wedge \partial^M(x \wedge y) \\
&= \delta^M \wedge \partial^M(x) \wedge y + (-1)^{|x|}\delta^M \wedge x \wedge \partial^M y \\
&= (\delta^V + \delta^{E-V}) \wedge \partial^M(x) \wedge y + (-1)^{|x|}(\delta^V + \delta^{E-V}) \wedge x \wedge \partial^M y \\
&= \delta^V \wedge \partial^M(x) \wedge y + \delta^{E-V} \wedge \partial^M(x) \wedge y + x \wedge \delta^V \wedge \partial^M y + x \wedge \delta^{E-V} \wedge \partial^M y \\
&= \lambda \cdot x \wedge y + \delta^{E-V} \wedge \partial^M(x) \wedge y + 0 + x \wedge \delta^{E-V} \wedge \partial^M y \\
&= \lambda \cdot x \wedge y + \delta^{E-V} \wedge \partial^M(x) \wedge y + x \wedge (\mu \cdot y + y')
\end{aligned}$$

where $y' \in \wedge^{r(M/V)}$ satisfies $x \wedge y' = 0$ (to be justified below). Assuming this, then the last expression is equal to the right-hand side in the lemma.

Now, the justifications. The first two equalities are by definition of Δ^M and the derivation property of ∂^M (equation (4.1)). The third and fourth equalities come from decomposing the sum $\delta^M = \delta^V + \delta^{E-V}$ and distributing over the wedge product. The fifth equality comes from the facts that

$$\delta^V \wedge \partial^M(x) = \delta^V \wedge \partial^V(x) (= \lambda \cdot x)$$
$$x \wedge \delta^V = 0$$

which follow since x is supported on bases of V. The last equality comes from the fact that

$$\delta^{M/V} \wedge \partial^{M/V} y = \mu \cdot y \quad \text{ in } \wedge M/V$$

implies

$$\delta^{E-V} \wedge \partial^M y = \mu \cdot y + y'$$

where y' is supported on subsets of $E - V$ which are dependent in M/V. However all the terms in y' will die in $\wedge M$ when one wedges with x, since x is supported on a set of bases of V. □

Corollary 13. *With the same conditions and notation of the previous lemma,*

(1) *if $\lambda = 0$, then*

$$\Delta^M(x \wedge y) = \mu \cdot x \wedge y.$$

(2) *if $r(V) = r(M) - 1$ and $y = \delta^{E-V}$ then $\mu = |E - V|$ and*

$$\Delta^M(x \wedge y) = (\lambda + |E - V|) \cdot x \wedge y.$$

Proof. To prove (1), note that $\lambda = 0$ implies x is in the 0-eigenspace of Δ^V, which implies that $\partial^V(x) = 0$. On the other hand, $\partial^V(x) = \partial^M(x)$ since x is supported on bases of V, so the extra term $\delta^{E-V} \wedge \partial^M(x) \wedge y$ in the right-hand side of the lemma dies.

To prove (2), the fact that $\Delta^{M/V}(y) = |E-V|\,y$ when $r(V) = r(M)-1$ and $y = \delta^{E-V}$ is easy to check. Also note that the extra term $\delta^{E-V} \wedge \partial^M(x) \wedge y$ dies because $\delta^{E-V} \wedge \delta^{E-V} = 0$. □

Given a flat V of M, we next define a map h_V

$$\begin{array}{rccccc} h_V: & (\Delta^V)_0 & \otimes & (\Delta^{M/V})_{|E-V|} & \to & C^M \\ & x & \otimes & y & \mapsto & x \wedge y \end{array}$$

where we recall that $(\Delta^M)_\lambda$ denotes the λ-eigenspace for the operator Δ^M. If we let $E(V)$ denote the image of h_V in C^M, then item (1) in Corollary 13 shows that $E(V)$ lies in the $|E-V|$-eigenspace for Δ^M. Since Δ^M is a self-adjoint operator, this implies that $E(V), E(V')$ are orthogonal whenever V, V' are flats of different cardinalities. Our goal will be to show more generally that $E(V), E(V')$ are always orthogonal, and that they give an orthogonal decomposition of the chain group C^M.

To achieve this goal, and for later purposes, it will be useful to know generators for the eigenspace $(\Delta^M)_{|E|}$. Given any maximal flag

$$F : \hat{0} \lessdot V_1 \lessdot \cdots \lessdot V_{r(M)-1} \lessdot M$$

in the geometric lattice $L(M)$, define $y_F \in \wedge M$ by

$$y_F := \left(\sum_{e \in V_1} v_e\right) \wedge \left(\sum_{e \in V_2 - V_1} v_e\right) \wedge \cdots \wedge \left(\sum_{e \in E - V_{r(M)-1}} v_e\right).$$

Using induction with item (2) in Corollary 13 shows that y_F lies in the eigenspace $(\Delta^M)_{|E|}$. We will not only show that these elements y_F span $(\Delta^M)_{|E|}$, but we will also relate them to the top cohomology group $\tilde{H}^{r(M)-2}(L(M))$ of the geometric lattice $L(M)$ (here we are abusing notation by using $L(M)$ to denote both the geometric lattice and the *order complex* of chains in the proper part of this lattice). Let $C^{r(M)-2}(L(M))$ denote the top cochain group over $\mathbb{R}$ for the order complex of the proper part of $L(M)$, and given a flag F in $L(M)$, let F^* denote the cochain which assigns the value 1 to F and 0 to all other flags.

Theorem 14. *The map $p : C^{r(M)-2}(L(M)) \to (\Delta^M)_{|E|}$ sending $F^* \mapsto y_F$ induces an isomorphism $p_* : \tilde{H}^{r(M)-2}(L(M)) \to (\Delta^M)_{|E|}$.*

Proof. First we must verify that p defined as above on the cochains actually gives a well-defined map on cohomology, i.e. p takes $(r(M)-1)$-coboundaries to 0. Since the $(r(M)-1)$-coboundaries are spanned by the coboundaries of F^* where F runs over all flags in $L(M)$ that miss only a single rank, it suffices to show that $p(\delta(F^*)) = 0$ for all such flags F. If F is such a flag, and it misses rank i, let V_{i-1}, V_{i+1} be its $i-1, i+1$-dimensional flats respectively. One can quickly check the computation $p(\delta(F^*)) = 0$ then amounts to a computation in the rank 2 geometric lattice $L(V_{i+1}/V_{i-1})$, where it can be easily verified to always work.

Since we have seen before that the dimensions of $\tilde{H}^{r(M)-2}(L(M)), (\Delta^M)_{|E|}$ are both equal to $|\mu_{L(M)}(\hat{0}, \hat{1})|$, it only remains to show that p_* is injective. To do this we will show that a particular basis for $\tilde{H}^{r(M)-2}(L(M))$ is mapped to a set of linearly independent elements in $(\Delta^M)_{|E|}$.

Given a base B of M having external activity 0, let $b_1 < \cdots < b_r$ be the elements of B in their ω-ordering, and let $F(B)$ be the complete flag

$$\overline{\{b_r\}} \lessdot \overline{\{b_r, b_{r-1}\}} \lessdot \cdots \lessdot \overline{\{b_r, b_{r-1}, \ldots, b_2, b_1\}}.$$

In [1, §7.6] it is shown that these flags are exactly the ones having decreasing label sets with respect to a certain EL-labelling of the Hasse diagram of $L(M)$. Combining this with [1, 7.6.3, 7.6.4, 7.7.2] and the fact that the Kronecker pairing gives an isomorphism $\tilde{H}_{r(M)-2}(L(M)) \cong \tilde{H}^{r(M)-2}(L(M))$, the cocycles $\{F(B)^*\}$ as B runs over all bases of M of external activity 0 give a basis for $\tilde{H}^{r(M)-2}(L(M))$.

Since p sends $F(B)^*$ to $y_{F(B)}$, it would suffice to show the following (cf. [1, proof of Theorem 7.8.4]): for any two bases B, B' of external activity 0, the coefficient of B in $y_{F(B)}$ is ± 1, while the coefficient of B' is 0. In proving this, we will use the following two observations, whose straightforward verifications are left to the reader:

(a) If B is a base of external activity 0, then any non-empty subset $B_0 \subseteq B$ is a base of external activity zero for the flat $\overline{B_0}$ (with respect to the ordering induced from ω).
(b) If $B = \{b_1, \ldots, b_r\}$ is a base of external activity 0, as above, and we set $B_i = \{b_r, b_{r-1}, \ldots, b_{i+1}, b_i\}$, then b_i is the ω-smallest element of $\overline{B_i} - \overline{B_{i+1}}$.

A typical term in the expansion of $y_{F(B)}$ corresponds to a sequence $(x_r, \ldots, x_1)$ with $x_j \in \overline{B_j} - \overline{B_{j+1}}$ for $1 \leq j \leq r$. We wish to show that $\{x_r, \ldots, x_1\}$ is a base of external activity 0 if and only if $x_j = b_j$ for all j. Using induction on j and fact (a) above, it suffices to show that given any $x \in \overline{B_j} - \overline{B_{j+1}}$, the set $X = \{b_r, b_{r-1}, \ldots, b_{j+1}, x\}$ is not a base of external activity 0 for $\overline{B_j}$. But this follows from fact (b) since $x \in ci_{\overline{B_j}}(X, b_j)$. □

We can now prove one of our main results. Recall that $E(V)$ is defined to be the image of h_V.

Theorem 15. *The map* $h := \bigoplus_V h_V$

$$\bigoplus_{flats\ V} (\Delta^V)_0 \otimes (\Delta^{M/V})_{|E-V|} \qquad \to \qquad C^M$$

is an isomorphism. Moreover

$$\bigoplus_{flats\ V\ of\ M} E(V) = C^M$$

is an orthogonal direct sum decomposition.

Proof. By Theorem 1, the dimensions of the two spaces are the same, and therefore it suffices to show that h is injective. For this we will show that

(a) for every flat V, the map

$$h_V : (\Delta^V)_0 \otimes (\Delta^{M/V})_{|E-V|} \to C^M$$

is injective, and
(b) for any two different flats V, V', the subspaces $E(V), E(V')$ are orthogonal.

Assertion (a) follows from the fact that h_V is a restriction of the injective map

$$\wedge^{|V|} V \otimes \wedge^{|E-V|} M/V \to \wedge^{|E|} M$$

given by $x \otimes y \mapsto x \wedge y$. To see that the latter map is injective, note that it maps a set of basis vectors for the domain into a subset of a basis for the range (the latter subset consisting of all v_B indexed by bases B of M which contain a base of V).

Assertion (b) is a bit trickier, and uses the Laplacian Δ^M in a fundamental way. Let V, V' be two different flats of M, and let $z = x \wedge y, z' = x' \wedge y'$ be elements of $E(V), E(V')$ respectively (so x, x' are in the eigenspaces of $(\Delta^V)_0, (\Delta^{V'})_0$ respectively, and y, y' are in the eigenspaces $(\Delta^{M/V})_{|E-V|}, (\Delta^{M/V'})_{|E-V'|}$ respectively). We wish to show $\langle z, z' \rangle = 0$.

If V, V' are flats of different cardinality, then we are done since then z, z' lie in different eigenspaces of the self-adjoint operator Δ^M.

If $|V| = |V'|$, then without loss of generality, by Theorem 14 we may assume that y, y' are of the form $y_F, y_{F'}$ for some complete flags F, F' in $L(M/V), L(M/V')$ respectively. Define a new matroid $\tilde{M}$ to be the *principal extension of M along the flat V by a new element p*; in the notation of [2], $\tilde{M} := M +_V p$. We recall a crucial property of principal extensions: any base B of M still forms a base in $\tilde{M}$. In this new matroid $\tilde{M}$, we have two flats $\tilde{V} := V \cup \{p\}$ and $\tilde{V'} = V'$ which have different cardinalities, and hence have $E(\tilde{V})$ orthogonal to $E(\tilde{V'})$ because they lie in different eigenspaces for $\Delta^{\tilde{M}}$. We can still think of x, x' as elements of $\wedge\tilde{M}$, and it is easy to see that they will still lie in the eigenspaces $(\Delta^{\tilde{V}})_0, (\Delta^{\tilde{V'}})_0$ because they are killed by $\partial^{\tilde{M}}$. Furthermore, when we think of y_F as an element of $\wedge^{\tilde{M}/\tilde{V}}$, it is still in the eigenspace $(\Delta^{\tilde{M}/\tilde{V}})_{|E-V|}$ because $\tilde{M}/\tilde{V} = M/V$. On the other hand, if we consider the chain of flats F' as a chain $\tilde{F'}$ of flats in $L(\tilde{M})$, it generates a new element $y_{\tilde{F'}}$. One can check that $y_{\tilde{F'}}$ differs from $y_{F'}$ only in some bases that contain p, and hence we have

$$\begin{aligned} \langle z, z' \rangle_M &= \langle x \wedge y_F, x' \wedge y_{F'} \rangle_M \\ &= \langle x \wedge y_F, x' \wedge y_{\tilde{F'}} \rangle_{\tilde{M}} \\ &= 0. \end{aligned}$$

The only tricky equality here is the second, which relies on the fact that $x \wedge y_F$ will not be supported on any bases that contain p. □

Question 16. *Is there any natural way to write down orthogonal projectors*

$$\pi_V : C^M \to E(V)$$

for each flat V of a matroid M?

In answering this question, we would consider choosing some arbitrary orthonormal basis for each $E(V)$ as not "natural" enough.

We can now state the main corollary to Theorem 15, interpreting the spectra of the Laplacians Δ_i^M. Recall that $(\Delta_i^M)_\lambda$ denotes the λ-eigenspace of Δ_i^M.

Corollary 17. *All Laplacians Δ_i^M for any matroid M have integral spectra. Furthermore, the generating function for these spectra is given explicitly by*

$$\sum_{\lambda \in \mathbb{R}, i \geq -1} \dim_{\mathbb{R}}(\Delta_i^M)_\lambda \, t^i q^\lambda = t^{-1} q^{|E|} \operatorname{Spec}_M(t, q^{-1}).$$

Proof. Beginning with the left hand side of the corollary, we perform a sequence of manipulations which are justified below.

$$\begin{aligned}
&\sum_{\lambda\in\mathbb{R}, i\geq -1} \dim_{\mathbb{R}}(\Delta_i^M)_\lambda t^i q^\lambda \\
&= \sum_{\text{flats } V \text{ of } M} t^{r(V)-1} \sum_{\lambda\in\mathbb{R}} \dim_{\mathbb{R}}(\Delta^V_{r(V)-1})_\lambda \, q^{\lambda+|E-V|} \\
&= \sum_{\text{flats } V \text{ of } M} t^{r(V)-1} \sum_{\text{flats } V'\subseteq V} \dim_{\mathbb{R}}(\Delta^{V'}_{r(V')-1})_0 \cdot \dim_{\mathbb{R}}(\Delta^{V/V'}_{r(V/V')-1})_{|V-V'|} \, q^{|V-V'|+|E-V|} \\
&= \sum_{\text{flats } V \text{ of } M} t^{r(V)-1} \sum_{\text{flats } V'\subseteq V} |\tilde{\chi}(IN(V'))| \cdot |\mu_{L(M)}(V',V)| \, q^{|E-V'|} \\
&= t^{-1} q^{|E|} \operatorname{Spec}_M(t, q^{-1})
\end{aligned}$$

The first equality above is justified by Proposition 5, the second equality by Theorem 15 and Corollary 13, the third equality by the same facts from [1] that were used in the proof of Theorem 1, and the last equality by equation (2.2). □

5. Symmetric examples

In [6], the analysis of the eigenspaces of the Laplacians for the chessboard complexes necessarily also entailed a description of these eigenspaces as representations for the symmetry groups present. Similarly, we can ask for such a description for the eigenspaces of the Laplacians of a matroid complex whenever there are nontrivial automorphisms of the matroid. This section is devoted to some examples of this nature.

We begin with a reduction. Combining Theorem 15 and Proposition 5 gives the isomorphism

(5.1)

$$\begin{aligned}
C_i^M &\cong \bigoplus_{\text{flats } V \text{ with } r(V)=i+1} \quad \bigoplus_{\text{flats } V'\subseteq V} (\Delta^{V'}_{r(V')-1})_0 \otimes (\Delta^{V/V'}_{r(V/V')-1})_{|V-V'|} \\
&\cong \bigoplus_{\text{flats } V \text{ with } r(V)=i+1} \quad \bigoplus_{\text{flats } V'\subseteq V} \tilde{H}_{r(V')-1}(IN(V')) \otimes \tilde{H}^{r(V/V')-2}(L(V/V'))
\end{aligned}$$

Let G be a group of automorphisms of the matroid M. Note that G permutes the chains of flats $V' \subseteq V$, and the stabilizer

$$Stab_G(V',V) := \{g \in G : gV' \subseteq V', gV \subseteq V\}$$

acts on $\tilde{H}_{r(V')-1}(IN(V'))$ and $\tilde{H}^{r(V/V')-2}(L(V/V'))$. Since the map h in Theorem 15 is easily checked to be G-equivariant, equation (5.1) leads immediately to the following result:

Theorem 18. *Let G be a group of automorphisms of M. Then as G-representations we have isomorphisms*

$$C_i^M \cong \bigoplus_{(V',V)} \operatorname{Ind}^G_{Stab_G(V',V)} \tilde{H}_{r(V')-1}(IN(V')) \otimes \tilde{H}^{r(V/V')-2}(L(V/V')) \tag{5.2}$$

as (V',V) runs over a set of representatives of all G-orbits of chains of flats $V' \subseteq V$ having $r(V) = i+1$.

In the statement of Theorem 18, $\mathrm{Ind}_{Stab_G(V',V)}^{G} W$ denotes the induction of a representation W of the subgroup $Stab_G(V', V)$ to the group G.

Note that all flats (V', V) in a single G-orbit will have V' of the same cardinality, and therefore their corresponding spaces $\tilde{H}_{r(V')-1}(IN(V')) \otimes \tilde{H}^{r(V/V')-2}(L(V/V'))$ will lie in the same eigenspace of Δ_i^M. This implies that the decomposition of G-representations in Theorem 18 is still finer than the decomposition into eigenspaces of Δ_i^M.

The previous result suggests that we concentrate attention on understanding the G-action on the homology of the matroid complex $IN(M)$. The actions on the cohomology of the geometric lattice $L(M)$ tend to have been studied earlier, or we can always use the isomorphism

$$\tilde{H}^{r(M)-2}(L(M)) \cong (\Delta^M)_{|E|} \cong (\Delta^{M^*})_0 \cong \tilde{H}_{r(M^*)-1}(IN(M^*)) \tag{5.3}$$

to reduce to the study of $IN(M^*)$ when this is convenient.

Notice that in general one can write down the G-action on $H_{r(M)-1}(IN(M))$ as a *virtual* character using the Euler characteristic:

$$\tilde{H}_{r(M)-1}(IN(M)) \cong \sum_{i \geq -1} (-1)^{r(M)-i-1} C_i^M. \tag{5.4}$$

However it is not always clear why this expression gives a genuine character rather than just a virtual one.

Example 1. Uniform matroids U(r,n)
The *uniform matroid* $M = U(r, n)$ of rank r on ground set $[n] := \{1, 2, \ldots, n\}$ has as independent sets all subsets of cardinality at most r. Flats, quotients, and duals of uniform matroids are all again uniform.

The symmetric group $G = S_n$ acts on $U(r, n)$ by permuting the ground set $[n]$, and the G-orbits of chains of flats $V' \subseteq V$ are determined by the cardinality of the sets V', V. In light of this discussion, it suffices to determine the S_n action on $IN(U(r, n))$. We can use equation (5.4), once we determine the S_n-action on C_i^M. The space C_i^M is the span of all oriented i-simplices $[j_0, \ldots, j_i]$. These i-simplices all lie in a single S_n-orbit. The stabilizer of a fixed oriented i-simplex is a subgroup isomorphic to $S_{i+1} \times S_{n-i-1}$, and we observe that the stabilizer acts on the 1-dimensional space spanned by this fixed oriented i-simplex by the sign character sgn of S_{i+1} tensored with the trivial character 1 of S_{n-i-1}. We conclude that

$$C_i^M \cong \mathrm{Ind}_{S_{i+1} \times S_{n-i-1}}^{S_n} sgn \otimes 1.$$

Using the Littlewood-Richardson rule [8, §4.9] along with equation (5.4) then shows that $IN(U(r, n))$ carries the irreducible S_n-representation indexed by the partition shape $(n - r, 1^r)$ (see [8, Chapter 2]).

Example 2. Finite projective spaces $\mathbb{P}_{\mathbb{F}_q}^{n-1}$
Let $\mathbb{F}_q$ be the finite field with q elements, for q a prime power, and $\mathbb{F}_q^n$ an n-dimensional space over $\mathbb{F}_q$. The projectivization of this space forms a matroid $M = \mathbb{P}_{\mathbb{F}_q}^{n-1}$ whose ground set is the set of lines through the origin in $\mathbb{F}_q^n$, and where a subset of lines is independent if and only if they contain a set of representing vectors which are linearly independent in the usual sense. Flats and quotients of finite projective spaces are again isomorphic to finite projective spaces. One might suspect from the terminology that dual matroids of finite projective spaces are

again isomorphic as matroids to finite projective spaces (for the dual space), but this is not true as shown by simple examples.

The finite general linear group $GL_n(\mathbb{F}_q)$ acts on $\mathbb{P}^{n-1}_{\mathbb{F}_q}$ in the obvious way and the G-orbits of chains of flats $V' \subseteq V$ are determined by the dimension of V', V as subspaces. In light of this discussion, it suffices to determine the $GL_n(\mathbb{F}_q)$ action on $\tilde{H}^{n-1}(L(\mathbb{P}^{n-1}_{\mathbb{F}_q}))$ and on $\tilde{H}_{n-1}(IN(\mathbb{P}^{n-1}_{\mathbb{F}_q}))$. The first task is relatively easy since $L(\mathbb{P}^{n-1}_{\mathbb{F}_q})$ is the *finite vector space lattice*, whose order complex is the *Tits building* for the usual BN-pair structure on $GL_n(\mathbb{F}_q)$. Here the $GL_n(\mathbb{F}_q)$-representation on the top (co-)homology group $\tilde{H}^{n-1}(L(\mathbb{P}^{n-1}_{\mathbb{F}_q}))$ is the well-known *Steinberg character* [9], an irreducible representation of $GL_n(\mathbb{F}_q)$ of dimension $q^{\binom{n}{2}}$.

To describe the $GL_n(\mathbb{F}_q)$-representation on $IN(\mathbb{P}^{n-1}_{\mathbb{F}_q})$, we can try to be more explicit in equation (5.4), by describing the $GL_n(\mathbb{F}_q)$-action on C_i^M. The space C_i^M is the span of all oriented i-simplices $[l_0, \ldots, l_i]$, for sets $\{l_0, \ldots, l_i\}$ of independent lines through the origin. Such i-simplicies all lie in a single $GL_n(\mathbb{F}_q)$-orbit, and the stabilizer of a fixed oriented i-simplex is a subgroup H isomorphic to the following subgroup of 2×2 block upper-triangular matrices:

$$\begin{pmatrix} A & B \\ 0 & C \end{pmatrix}$$

Here A must be an $(i+1) \times (i+1)$ *monomial* matrix, i.e. it has exactly one non-zero entry in each row and column, B is an arbitrary $(i+1) \times (n-i-1)$ matrix, and C is an invertible $(n-i-1) \times (n-i-1)$ matrix. Elements of the stabilizer H act on the oriented simplex which they fix by the sign of the permutation associated to the monomial matrix A, so call this character χ. We conclude that $C_i^M \cong \mathrm{Ind}_H^{GL_n(\mathbb{F}_q)} \chi$, and equation (5.4) gives an expression for $\tilde{H}_{n-1}(IN(\mathbb{P}^{n-1}_{\mathbb{F}_q}))$ as a virtual character of dimension

$$\sum_{i \geq -1} (-1)^{n-1-i} \frac{[n]_q [n-1]_q \cdots [n-i]_q}{(i+1)!} q^{\binom{i+1}{2}} \tag{5.5}$$

where $[n]_q := 1 + q + q^2 + \cdots + q^{n-1}$. We do not see how to simplify the expression (5.4) for $\tilde{H}_{n-1}(IN(\mathbb{P}^{n-1}_{\mathbb{F}_q}))$ as a virtual character.

Example 3. The graphic matroid for a complete graph

Let M_n^A be the graphic matroid associated to the complete graph on vertex set $[n]$. In other words, M_n^A has as ground set E the set of all possible edges $\{i, j\}$ with $1 \leq i < j \leq n$, and a set of edges is independent if it contains no cycles.

The symmetric group S_n acts on M_n^A by permuting the vertex set $[n]$, and hence also permuting the edges $\{i, j\}$. Flats in M_n^A correspond to partitions $\pi = B_1|B_2|\cdots|B_r$ of $[n]$ into blocks B_i, and one can check that if V is the flat corresponding to the partition π then as a matroid $V \cong \bigoplus_i M^A_{|B_i|}$. The lattice $L(M_n^A)$ is the well-known lattice of set partitions of $[n]$, usually denoted Π_n. Furthermore, if V' is a flat contained in V, then V' corresponds to a partition π' which refines π, and $L(V/V')$ is isomorphic to a product of lattices of the form Π_{n_i} for some $n_i < n$.

In light of this discussion, it suffices to determine the S_n-action on $\tilde{H}^{n-3}(\Pi_n)$ and on $\tilde{H}_{n-2}(IN(M_n^A))$. The first task is relatively easy since the S_n-action on $\tilde{H}^{n-3}(\Pi_n)$ has been well-studied by many authors, see e.g. [10, §7]. It is known

that as an S_n-representation,

$$\tilde{H}^{n-3}(\Pi_n) \cong sgn \otimes \mathrm{Ind}_{C_n}^{S_n} \zeta$$

where sgn is the 1-dimensional sign representation, C_n is the cyclic group in S_n generated by an n-cycle, and ζ is the 1-dimensional representation of C_n which sends a generator to a primitive n^{th} root of unity [10, Theorem 7.3]. It is also known that if we let S_{n-1} denote the subgroup of S_n which fixes n, then

$$\mathrm{Res}_{S_{n-1}}^{S_n} \tilde{H}^{n-3}(\Pi_n) \cong \mathbb{R}[S_{n-1}]$$

i.e. the restriction to S_{n-1} is isomorphic to the regular representation of S_{n-1} [10, Corollary 7.6]. Consequently, this representation has dimension $(n-1)!$.

The S_n-action on $\tilde{H}_{n-2}(IN(M_n^A))$ has only been studied more recently in [7]. The description of the S_n-representation is not known in general, although equation (5.4) can be used to decompose it into S_n-irreducibles via a character computation when n is small. On the other hand, one can construct an explicit basis for the homology which is permuted by the action of S_{n-1} and which shows the following:

Theorem 19. [7, Theorem 3.6.3] *As an S_{n-1}-representation $\tilde{H}_{n-2}(IN(M_n^A))$ is isomorphic to the permutation action on forests of edge-rooted trees on $[n-1]$ tensored with the sign character sgn.*

Here a forest of edge-rooted trees is a graph on $[n-1]$ in which every connected component is a tree, along with a specified edge of each tree (so in particular, every vertex in $[n-1]$ lies in a tree having at least two vertices). We give a brief proof of this result here, relying on results of [1].

First, choose the ordering ω of the edges in the complete graph as follows: an edge $\{i,j\}$ with $i<j$ is earlier in ω than $\{i',j'\}$ if $j>j'$ or if $j=j'$ and $i>i'$. In other words

$$\{n-1,n\} <_\omega \cdots <_\omega \{1,n\} <_\omega \{n-2,n-1\} <_\omega \cdots <_\omega \{2,3\} <_\omega \{1,3\} <_\omega \{1,2\}.$$

Bases of M_n^A correspond to spanning trees T on $[n]$, and it is straightforward to check the following characterization of the bases of internal activity 0:

Lemma 20. *A spanning tree T on $[n]$ has internal activity 0 if and only if for every tree of the forest $T|_{[n-1]}$, the vertex connecting this tree to n does not have the maximum label in the component.*

Here $T|_{[n-1]}$ denotes the restriction of the tree T to its edges which only contain vertices in $[n-1]$.

If we let $T_1, \ldots, T_r$ denote the trees of internal activity 0, then [1, Theorem 7.7.2] implies that there is a basis of cycles $z_1, \ldots, z_r$ for $\tilde{H}_{n-2}(IN(M_n^A))$ uniquely determined as homology classes by the condition that the coefficient of T_i in z_j is 1 if $i=j$ and 0 else. These cycles $\{z_i\}$ turn out *not* to be permuted by the S_{n-1}-action, however they are upper-triangularly related to a set of cycles which are so permuted, and which we now define.

Given a forest F of edge-rooted trees on $[n-1]$, let $F_1, \ldots, F_s$ be its component edge-rooted trees, let the distinguished edge in tree F_r be $\{i_r, j_r\}$, and for any other edge $e \neq \{i_r, j_r\}$ in F_r let $\psi(e)$ be the vertex in this edge which lies farthest from the root edge. See Figure 2 for some examples.

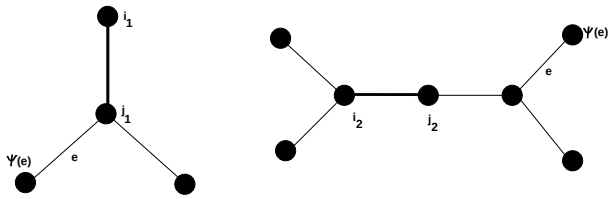

FIGURE 2. Example of an edge rooted-forest along with some non-root edges e and their corresponding vertices $\psi(e)$.

Define z'_F to be the following product in the exterior face ring $\wedge M_n^A$:

$$z'_F := \bigwedge_{r=1}^{s} \left(\partial(v_{i_r j_r} \wedge v_{i_r n} \wedge v_{j_r n}) \wedge \bigwedge_{\text{edges } e \neq \{i_r, j_r\} \text{ in } F_r} \partial(v_e \wedge v_{\psi(e),n}) \right)$$

We next describe a bijection ϕ between the forests of edge-rooted trees on $[n-1]$ and the trees of internal activity 0 on $[n]$. Given such a forest F let $\phi(F)$ be the tree on $[n]$ obtained as follows: for each edge-rooted tree F_r of F, assume that the labelling i_r, j_r of the distinguished edge has been chosen so that i_r is farther than j_r is from the maximum labelled vertex of F_r. Then $\phi(F)$ is obtained by adding all the edges $\{i_r, n\}$ for $1 \leq r \leq s$ to F. It is easy to see that $\phi(F)$ fits the description of Lemma 20, and so has internal activity 0. The inverse map ϕ^{-1} starts with a tree T of internal activity 0, removes all edges involving the vertex n leaving a forest on $[n-1]$, and for each vertex i that used to be connected to n it distinguishes the edge containing i in the forest which is the first edge on the path toward the maximum labelled vertex of i's component tree.

We now come to the crucial point. Given two forests F_i, F_j of edge-rooted trees, having corresponding trees T_i, T_j of internal activity zero, if the cycle z'_{F_i} contains the tree T_j, then one can easily check that:

(a) as a forest, F_i contains all the edges in F_j, and furthermore
(b) if F_i and F_j have exactly the same set of edges as forests, then in each of their component trees, the distinguished edge for F_j lies on the shortest path connecting the distinguished edge for F_i to the maximum labelled vertex in that component.

Consequently, if we linearly order the forests of edge-rooted trees by any linear extension of the partial order expressed in items (a), (b) above, we have the following property: when one restricts the matrix expressing the cycles z'_{F_i}'s in terms of trees on $[n]$ to the trees of internal activity 0, the result is a square upper-triangular matrix U with ± 1's on the diagonal. By [1, Theorem 7.7.2], this implies the cycles $\{U^{-1} z'_F\}$ must induce the exact same homology classes as the basis $\{z_i\}$ mentioned earlier for $\tilde{H}_{n-2}(IN(M_n^A))$. The next corollary is then immediate:

Corollary 21. *The cycles $\{z'_F\}$ as F ranges over all forests of edge-rooted trees form a basis for $\tilde{H}_{n-2}(IN(M_n^A))$.*

It is easy to see that permutations in S_{n-1} permute the cycles $\{z'_F\}$ (up to sign) in the same way that they permute the forests of edge-rooted trees. Furthermore, it is not hard to check that if F is a forest of edge-rooted trees which is fixed by some permutation π in S_{n-1}, then we also have $\pi(z'_F) = sgn(\pi) \cdot z'_F$. This proves Theorem 19.

6. Acknowledgments

The authors would like to thank Paul Edelman for helpful comments.

References

[1] A. Björner, *Homology of matroids*, in Matroid Applications (ed. N. White), Encyclopedia of Mathematics and Its Applications, **40**, Cambridge Univ. Press 1992.

[2] T. Brylawski, *Constructions*, in Theory of Matroids (ed. N. White), Encyclopedia of Mathematics and Its Applications, **26**, Cambridge Univ. Press 1986.

[3] T. Brylawski and J. G. Oxley, *The Tutte polynomial and its applications*, in Matroid Applications (ed. N. White), Encyclopedia of Mathematics and Its Applications, **40**, Cambridge Univ. Press 1992.

[4] A. Björner and G. Kalai, *On f-vectors and homology*, in Combinatorial Mathematics: Proc. of 3rd International Conference (New York, 1985), N.Y. Acad. Sci.: New York, 1989.

[5] J. Friedman, *Computing Betti numbers via combinatorial Laplacians*, in Proc. 28th Annual ACM Symposium on the Theory of Computing, ACM: New York, 1996.

[6] J. Friedman and P. Hanlon, *On the Betti numbers of chessboard complexes*, preprint 1996.

[7] W. Kook, *Categories of acyclic graphs and automorphisms of free groups*, Ph.D. thesis (G. Carlsson, advisor), Stanford Univ. 1996.

[8] B. Sagan, The symmetric group: representations, combinatorial algorithms, and symmetric functions, Wadsworth & Brooks/Cole: Pacific Grove, 1991.

[9] L. Solomon, *The Steinberg character of a finite group with a BN-pair*, in Theory of Finite Groups (eds. R. Brauer and C. H. Sah), Benjamin: New York, 1969, 213-221.

[10] R. Stanley, *Some aspects of groups acting on finite posets*, J. Combin. Theory Ser. A, **32** (1982) 132–161.

E-mail address: kook@math.umn.edu

E-mail address: reiner@math.umn.edu

E-mail address: stanton@math.umn.edu

School of Mathematics, University of Minnesota, Minneapolis, MN 55455