A $GL_n(q)$ ANALOGUE OF THE PARTITION LATTICE

PHIL HANLON, PATRICIA HERSH, AND JOHN SHARESHIAN

Abstract. This paper introduces a $GL_n(q)$ analogue for the partition lattice, namely the lattice of partial direct sum decompositions of a finite vector space, denoted $PD_n(q)$. $PD_n(q)$ is proven to be a homotopically Cohen-Macaulay poset, implying its homology is concentrated in top degree. There is an action of $GL_n(q)$ on $PD_n(q)$ giving rise to a $GL_n(q)$-representation on top homology. In analogy to the S_n-representation on the top homology of the partition lattice, this $GL_n(q)$-representation is proven to be induced from a linear character of the normalizer of a Coxeter torus which is trivial on the torus and faithful on a complement. The lattice $\Pi_{\leq n}$ of partial partitions of a finite set is also introduced and shown to be a collapsible, supersolvable lattice for which $PD_n(q)$ is the q-analogue. $PD_n(q)$ is covered by copies of $\Pi_{\leq n}$ indexed by the bases for a finite projective space; discrete Morse theory is used to construct $PD_n(q)$ by sequentially attaching these copies of $\Pi_{\leq n}$ in a way that preserves the Cohen-Macaulay property at each step.

1. Introduction

One important motivation for the study of poset homology is the wealth of interesting group representations that arise in this context. If P is a graded poset and G is a group of automorphisms of P, then G acts on the (reduced) homology, $\widetilde{H}_*(P_S)$, of any rank-selected subposet of P. Whereas the representation of G determined by its action on each chain space $C_i(P_S)$ is fairly straightforward (being a permutation representation), there are numerous examples in which the representation of G on $\widetilde{H}_*(P_S)$ has deep and interesting properties. This approach is particularly fruitful when the poset P is Cohen-Macaulay. In that case, $\widetilde{H}_*(P_S)$ is concentrated in top dimension $t(S)$ for each S, and the character value of $g \in G$ acting on $\widetilde{H}_{t(S)}(P_S)$ can be determined via a combinatorial Möbius function computation.

The case in which $P = \Pi_n$, the lattice of partitions of $\{1, \ldots, n\}$, and $G = S_n$, the symmetric group on n letters, provides one of the richest

The second author was supported by an NSF Postdoctoral Research Fellowship and NSF grant 0500638. The third author was supported by NSF grant 0300483.

instances of this construction. Building on the character value computation by Hanlon [Ha1], Stanley initiated the study of the representation of the symmetric group S_n on the top homology group $\widetilde{H}_{n-3}(\Pi_n)$ of the full partition lattice. In [St3], Stanley showed that the representation of S_n on this group is equivalent (up to a sign twist) to Lie_n, the representation of S_n on the multilinear part of the free Lie algebra with n generators. Barcelo [Ba] gave a constructive proof of this fact by displaying an S_n-equivariant isomorphism between $H_{n-3}(\Pi_n)$ and $sgn \otimes Lie_n$, where sgn is the sign representation of S_n. In later work, Wachs provided a basis-free proof of this fact in [Wac]. Subsequently, a number of elegant results have been proved about the homology of subposets of Π_n, including rank selections of Π_n and Cohen-Macaulay subposets consisting of elements chosen by the numerological properties of their block sizes (see [Ha2], [Su1], [Su2], [Su3], [Su4], [HH], [CHR], [HaW], [SW]). It is known (see for example [St3]) that if $\sigma \in S_n$ is an n-cycle then Lie_n is the representation of S_n induced from any faithful linear character of the group $\langle \sigma \rangle$ generated by σ. Note that $\langle \sigma \rangle = C_{S_n}(\sigma)$.

Our aim is to suggest an analogous program for $GL_n(q)$. To this end, we introduce a poset called $PD_n(q)$ to play the role of the partition lattice. The elements of $PD_n(q)$ are obtained as follows: choose a subspace W of the n-dimensional vector space V over the finite field F_q; now take a collection of subspaces which together span W and whose sum is direct. Now order elements by $u = \{U_1, \ldots, U_k\} < w = \{W_1, \ldots, W_l\}$ if each component U_i of u is a subspace of some component W_j of W. For example, let $\{e_1, \ldots, e_n\}$ be a basis for V, let U be spanned by $e_1 + e_2 + e_4$, let W be spanned by $e_1 + e_2$ and e_4, let X be spanned by $e_2 + e_3$, let Y be spanned by $e_1 - e_2$ and e_3 and let Z be spanned by e_1, e_2, e_3 and e_4. Then, in $PD_n(q)$, we have $\{U\} < \{W\} < \{W, Y\} < \{Z\}$, while $\{X\}$ and $\{W, Y\}$ are not related.

Our first main theorem is the following one, which lays the groundwork for a study of the $GL_n(q)$-representation on homology by showing that homology is again concentrated in top degree. As noted above, the following will also imply, via a result of Munkres [Mu, Cor. 6.6] (see also [St1], [Wal]), that all rank-selected subposets also have homology concentrated in top degree.

Theorem 1.1. *$PD_n(q)$ is a Cohen-Macaulay poset.*

We then determine the representation of $GL_n(q)$ on the unique nontrivial homology group $\widetilde{H}_{2n-3}(PD_n(q))$ by computing its character. As we will discuss presently, it turns out that this representation bears a striking resemblance to that of S_n on the homology of the partition lattice. Indeed, this is why we have chosen to study $PD_n(q)$ as a $GL_n(q)$-analogue of the partition lattice rather than the lattice of direct sum decompositions of a finite vector space, which was proven to be Cohen-Macaulay by Welker in [We]. The (not yet well understood) representation of $GL_n(q)$ on the homology of the lattice of full direct sum decompositions does not seem to resemble the representation of S_n on the homology of Π_n in any obvious way.

Recall (or see for example [Ca, Section 3.3]) that for each partition $\lambda = (\lambda_1, \dots, \lambda_r)$ of n, $GL_n(q)$ has a unique conjugacy class of abelian subgroups isomorphic to $\prod_{i=1}^r \mathbb{Z}_{q^{\lambda_i}-1}$. These groups (as λ runs through all partitions of n) are called the *maximal tori* of $GL_n(q)$. We shall be interested in the maximal tori which are cyclic of order $q^n - 1$ (and therefore correspond to the partition (n)), which are called the *Coxeter tori* of $GL_n(q)$.

We review the basic properties of the Coxeter tori and their normalizers before stating our second main result. Identify the vector space $\mathbb{F}_q^n$ with the field $\mathbb{F}_{q^n}$. Let α generate the multiplicative group $\mathbb{F}_{q^n}^\times$, and consider the map from $\mathbb{F}_{q^n}$ to itself

$$c_\alpha : \beta \mapsto \beta\alpha.$$

This map is a nonsingular linear transformation of order $q^n - 1$ (as it permutes the elements of $\mathbb{F}_{q^n}^\times$ cyclically). The Coxeter tori in $GL_n(q)$ are those subgroups which are generated by conjugates of c_α. Let T be the Coxeter torus generated by c_α. Now consider the Frobenius automorphism on $\mathbb{F}_{q^n}$,

$$f : \beta \mapsto \beta^q.$$

We have

$$f^{-1} c_\alpha f = c_\alpha^q,$$

so f normalizes T. Moreover, by the Normal Basis Theorem, f cyclically permutes an $\mathbb{F}_q$-basis for $\mathbb{F}_{q^n}$. Thus f is conjugate to a permutation matrix corresponding to an n-cycle in S_n. Let N be the normalizer of T in $GL_n(q)$. We have seen that $\langle f \rangle T \leq N$, and it turns out that $N = \langle f \rangle T$.

We can now state our result on the representation of $GL_n(q)$ on $\widetilde{H}_{2n-3}(PD_n(q))$.

Theorem 1.2. *Let $\theta_n : N \to \mathbb{C}^\times$ be the unique representation such that $\theta_n(T) = 1$ and $\theta_n(f) = e^{2\pi i/n}$. Let Θ_n be the representation of $GL_n(q)$ obtained by inducing the representation θ_n of N up to $GL_n(q)$. Then the representation of $GL_n(q)$ on $\widetilde{H}_{2n-3}(PD_n(q))$ is equivalent to Θ_n.*

An immediate corollary is the following Möbius function expression, obtained by computing the dimension of the representation Θ_n.

Corollary 1.3. *If q is a prime power, then* $\mu_{PD_n(q)}(\hat{0},\hat{1}) = -\frac{1}{n} \cdot q^{\binom{n}{2}} \prod_{i=1}^{n-1}(q^i - 1)$.

The similarity of Theorem 1.2 and Stanley's result on the partition lattice, and thus the justification for studying partial decompositions rather than full decompositions as mentioned above, are hopefully now apparent. This similarity is reinforced by the fact (again, see [Ca]) that if X is a maximal torus in $GL_n(q)$ which corresponds to a partition of shape λ and σ is an element of cycle shape λ in S_n then $N_{GL_n(q)}(X)/X \cong C_{S_n}(\sigma)$. Here, we have "transferred" a result about a representation induced from a faithful linear character of the centralizer of an element of shape (n) to a result about a representation induced from a linear character of the normalizer of a maximal torus of "shape" (n) whose kernel contains the torus, which is faithful on a complement to the torus.

We now describe briefly our proofs of Theorems 1.1 and 1.2. Theorem 1.1 is proven using a generalized lexicographic discrete Morse function; several rounds of critical cell cancellation in Section 8 allow us to transform this into a Morse function with only top-dimensional critical cells, as needed to deduce the homotopy type. The facet order is based on a decomposition of $PD_n(q)$ into subcomplexes indexed by the bases for the matroid of independent lines in a finite vector space. In Section 4 we show that each such subcomplex is isomorphic to the order complex of the lattice $\Pi_{\leq n}$ of partial partitions of $[n]$. This partial partition lattice is introduced and shown to be EL-shellable (in fact supersolvable) and collapsible in Section 3. In Section 11, we show that $PD_n(q)$ is in fact a q-analogue of the lattice of partial partitions, in the sense that statistics such as characteristic polynomial and zeta polynomial for $PD_n(q)$ specialize to those of $\Pi_{\leq n}$ as q approaches 1.

We build the order complex of $PD_n(q)$ by attaching these subcomplexes sequentially. Properties of the circuits in the given matroid allow us to analyze the overlap between the subcomplexes via discrete Morse theory. In analogy with the standard way of shelling a matroid independence complex (see [Bj2]), the only subcomplexes whose attachment may change the homotopy type are those indexed by bases with internal activity zero, but now we attach collapsible subcomplexes at each step, rather than just facets. With considerable effort, we show that when subcomplexes corresponding to arbitrary bases are attached, the homotopy type either remains unchanged or is adjusted by taking the wedge of the complex built so far with some spheres of highest possible dimension. Specifically, we use a property of the matroid of lines in a finite vector space which we call Property R, and which is further developed in Section 12. To deduce the Cohen-Macaulay property, the homotopy type result is generalized in Section 9 to arbitrary poset intervals, using ideas related to matroid contraction.

We prove Theorem 1.2 in Section 10 using induction on n. Recall (see for example [St3]) that the lattice $B_n(q)$ of subspaces of $\mathbb{F}_q^n$ is Cohen-Macaulay of dimension $n-2$, and the representation of $GL_n(q)$ on $\widetilde{H}_{n-2}(B_n(q))$ is the (irreducible) Steinberg representation, which will be denoted by St. We use the Whitney homology of $PD_n(q)$ to express St as the sum of virtual $GL_n(q)$-module induced from modules of stabilizers of elements $X = \{X_1, \ldots, X_k\} \in PD_n(q)$ such that $\langle X_1, \ldots, X_k \rangle = \mathbb{F}_q^n$. In fact, we choose one such element from each $GL_n(q)$-orbit. The virtual module associated to $\hat{1} = \mathbb{F}_q^n$ (with stabilizer $GL_n(q)$) is $-\widetilde{H}_{2n-3}(PD_n(q))$, and our inductive hypothesis is used to show that the remaining modules can be induced from linear characters of normalizers of maximal tori. This allows us to invoke a theorem of Srinivasan, thereby reducing our task to showing that some complicated but explicit identities hold in the character ring of the Weyl group S_n.

The paper is organized into three main parts. In Sections 2-6 we discuss general structure of $PD_n(q)$ and background. In Sections 7-10 we use discrete Morse theory to compute the homotopy type of $PD_n(q)$. In Section 11 we determine the homology representation, providing the second main result. In addition, in Section 12 we prove that $PD_n(q)$ is the q-analogue of the partial partition lattice, and in Section 13 we suggest possible generalizations to partial decomposition posets resulting from other matroids.

2. Background on Cohen-Macaulay and shellable posets

In this section we briefly review terminology and concepts, many of which appear in greater detail in [Bj1] and [St2]. Let P be a finite graded poset with a unique minimum element $\hat{0}$ and a unique maximum element $\hat{1}$. Denote by rk(u) the rank of any $u \in P$. We say P has rank n if $rk(\hat{1}) = n$. The **order complex** of $P \setminus \{\hat{0}, \hat{1}\}$, denoted $\Delta(P \setminus \{\hat{0}, \hat{1}\})$, is the simplicial complex whose i-faces are the chains $\hat{0} < x_0 < x_1 < \cdots < x_i < \hat{1}$ of $i+1$ comparable poset elements. Sometimes for convenience we will speak of the order complex of P, but in this case we are always referring to the order complex of $P \setminus \{\hat{0}, \hat{1}\}$. The reduced Euler characteristic of $\Delta(P \setminus \{\hat{0}, \hat{1}\})$ gives a way of computing the Möbius function (as defined in [St2]) by the formula

$$\mu_P(\hat{0},\hat{1}) = \tilde{\chi}(\Delta(P \setminus \{\hat{0},\hat{1}\})).$$

A simplicial complex is Δ is **pure** if all facets (maximal faces) of Δ have the same dimension. The complex Δ is **Cohen-Macaulay** if it has integral homology concentrated in top degree, and the same holds for the link of each of its faces. A poset is **Cohen-Macaulay** if its order complex is Cohen-Macaulay, which can be proven by showing that the order complex of each interval (x, y), denoted $\Delta(x, y)$, has homology concentrated in top degree. Recall that (x, y) consists of those poset elements z satisfying $x < z < y$. A rank-selected subposet P^S of a graded poset P of rank n is the subposet consisting of the elements of P whose rank belongs to $S \subseteq \{1, \ldots, n-1\}$. As mentioned earlier, it was proven in [Mu] that the rank-selected subposets of a Cohen-Macaulay poset are themselves Cohen-Macaulay.

One way of proving that a poset is Cohen-Macaulay is by showing that it is shellable. Recall that a **shelling** on a pure simplicial complex is an ordering $F_1, \ldots, F_k$ of its facets such that each subcomplex $F_j \cap (\cup_{i<j} F_i)$ is pure of codimension one. A poset is **shellable** if its order complex has a shelling. Notice that in a shelling each $F_j \cap (\cup i < jF_i)$ is either collapsible (because it includes one or more cone points), or it is the entire boundary of F_j. This implies that a shellable complex is homotopy equivalent to a wedge of top-dimensional spheres, and in fact that it is Cohen-Macaulay. Letting $F_1 \cup \cdots \cup F_j$ denote the subcomplex comprised of faces contained in one or more facets $F_1, \ldots, F_j$, refer to the process of transforming $F_1 \cup \cdots \cup F_{j-1}$ into $F_1 \cup \cdots \cup F_j$ as the **attachment** of F_j. Notice that the Möbius function $\mu_P(\hat{0},\hat{1})$ of a shellable, graded poset P equals the number of facets attaching along their entire boundary multiplied by $(-1)^{\mathrm{rk(P)}}$. Now we turn to a convenient way of constructing a shelling for $\Delta(P)$, namely the theory of lexicographic shellability, due to Björner and Wachs (see [Bj2], [BW]).

An **edge labelling** of P is a function λ from the set of pairs (x, y) such that y covers x in P to a totally ordered label set Q (or equivalently, to the integers). A chain $C : x_0 < x_1 < \ldots < x_k$ in P is called **saturated** if $x_i \le z \le x_{i+1}$ implies $z = x_i$ or $z = x_{i+1}$, i.e. x_{i+1} **covers** x_i. Each saturated chain $\hat{0} = x_0 < \cdots < x_r = \hat{1}$ is labelled by an r-tuple of labels $(\lambda(x_0, x_1), \ldots, \lambda(x_{r-1}, x_r))$ called its *label sequence.* The **content** of a label sequence is the multiset of labels appearing in it.

Lexicographically ordering label sequences gives a partial order on the saturated chains in P, i.e. on the facets of $\Delta(P \setminus \{\hat{0},\hat{1}\})$. We will be interested in **linear extensions** of this partial order on facets, i.e. total orders which are compatible with the partial order. The labelling λ is called an **EL-labelling** if whenever $x < y$ in P, there exists a unique saturated chain $C : x = x_0 < \ldots < x_m = y$ from x to y such that $\lambda(x_{i-1}, x_i) \leq \lambda(x_i, x_{i+1})$ for $1 \leq i \leq m$ (such a chain is called *rising*) and for any other saturated chain D from x to y, we have $\lambda(C) < \lambda(D)$. The poset P is **EL-shellable** if P admits an EL-labelling. Next we list a few key properties of EL-shellable posets.

- If P admits an EL-labelling λ then every linear extension of the lexicographic order on label sequences will be a shelling order on the facets of $\Delta(P \setminus \{\hat{0}, \hat{1}\})$.
- If we attach facets sequentially according to such a total order, then any facet whose label sequence includes an ascent, i.e. a pair of consecutive labels $\lambda(x_{i-1}, x_i) \leq \lambda(x_i, x_{i+1})$, will attach without changing the homotopy type because each ascent gives rise to a cone point in the intersection of the facet with the union of earlier facets.
- Any facet whose label sequence has no such ascents will attach along its entire boundary, increasing by one the number of spheres in the wedge of spheres. (Such saturated chains are called **decreasing chains**, and counting such chains can be an efficient way to compute Möbius functions.)

The partition lattice is a well-known example of an EL-shellable poset. Recall that the **partition lattice**, denoted Π_n, consists of the partitions of $[n] := \{1, \ldots, n\}$ into unordered components, with these partitions ordered by refinement. That is, we have $u < v$ if and only if each component of v is a union of components of u. It is shown in [Bj2] that the following labelling is an EL-labelling on Π_n. Label the covering relation $u \prec v$ which merges a pair of components B_1, B_2 of u into a single component in v with the integer $\lambda(u, v) = \max(\min B_1, \min B_2)$. For example, if $u = 1,3|2,5,8|4,6,7,9,10$ and $v = 1,3|2,4,5,6,7,8,9,10$ then $\lambda(u, v) = 4$.

A **modular element** in a graded lattice is an element u which satisfies $\mathrm{rk(u)} + \mathrm{rk(v)} = \mathrm{rk(u \vee v)} + \mathrm{rk(u \wedge v)}$ for every v, and a lattice which has a saturated chain consisting entirely of modular elements is **supersolvable**.

Remark 2.1. *$PD_n(q)$ is a graded lattice, but if $n > 2$ then $PD_n(q)$ contains no modular element other than $\hat{0}$ and $\hat{1}$. In particular, $PD_n(q)$ is not supersolvable in general. Since these facts are not essential to this paper, the proofs are omitted. It is not known whether or not $PD_n(q)$ is shellable, despite efforts to find a shelling.*

We will show that $PD_n(q)$ is Cohen-Macaulay, using the fact that it is a union of subposets that are EL-shellable, and which overlap in a way that we can analyze by discrete Morse theory. In the next section we introduce these subposets and provide an EL-labelling that is related to the one given above for the partition lattice.

3. The lattice of partial partitions of a finite set

In this section we introduce the lattice $\Pi_{\leq n}$ of partial partitions of a finite set $[n]$ and confirm a few basic properties. In the next section we will show that $PD_n(q)$ is a union of subposets, each of which is isomorphic to $\Pi_{\leq n}$.

Definition 3.1. *Let $\Pi_{\leq n}$ be the poset of partitions of subsets of $[n]$, ordered by refinement. That is, we have $\sigma \leq \rho$ whenever each component of σ is contained in a component of ρ.*

We write an arbitrary element $\pi \in \Pi_{\leq n}$ as $\pi = P_1|\ldots|P_k$, where the components P_i are pairwise disjoint subsets of $[n]$. In this context, the usual partition lattice Π_n is the sublattice of $\Pi_{\leq n}$ consisting of those elements which are greater than or equal to $\{1\}|\ldots|\{n\}$. We say that $\rho \in \Pi_{\leq n}$ is obtained from π by a *creation step* if there is some $j \in [n] \setminus \bigcup_{i=1}^{k} P_i$ such that the components of ρ are those of π along with $\{j\}$. We say that ρ is obtained from π by a *merge step* if there exist (distinct) $i, j \in [k]$ such that the components of ρ are $P_i \cup P_j$ along with P_l for all $l \in [k] \setminus \{i, j\}$. Note that for $\pi, \rho \in \Pi_{\leq n}$, ρ covers π if and only if ρ is obtained from π by either a creation step or a merge step. It follows that $\Pi_{\leq n}$ is graded, with every maximal chain having length $2n - 1$.

One may obtain a bijection from $\Pi_{\leq n}$ to the Dowling lattice $Q_n[1]$ (see [Do]) by mapping $\sigma \in \Pi_{\leq n}$ to the partition whose restriction to the subset of $[n]$ whose elements appear in σ is σ itself and whose 0-block contains the remaining elements of $[n]$. However, this map is not order preserving, and it is straightforward to see that $\Pi_{\leq n}$ is not isomorphic to $Q_n[1]$.

Proposition 3.2. $\Pi_{\leq n}$ *is a lattice.*

PROOF. Given two partial partitions σ, τ, we describe their unique meet and join. Obtain $\overline{\sigma}$ from σ and $\overline{\tau}$ from τ by adding to σ (resp. τ) blocks of size one, consisting of each set element appearing in τ but not in σ (resp. σ but not in τ). Then $\sigma \vee \tau$ is the join $\overline{\sigma} \vee \overline{\tau}$ in the partition lattice on the set partitioned by $\overline{\sigma}$ and by $\overline{\tau}$.

Similarly, the meet is obtained by restricting u, v to the intersection of their supports, then taking the meet on the partition lattice with this support. $\square$

Next we define an edge labelling λ on $\Pi_{\leq n}$ and show it is an EL-labelling, a result that will be essential to our upcoming analysis of $PD_n(q)$. Let Q be the set $\{c, m\} \times [n]$ with the order given by $(c, i) < (m, j)$ for all i, j and, for $x \in \{c, m\}$, $(x, i) \leq (x, j)$ if and only if $i \leq j$. Define the edge labelling λ by

$$\lambda(\pi, \rho) := \begin{cases} (c, i) & \text{if } \rho \text{ is obtained from } \pi \text{ by a creation step which adds} \\ & \text{component } \{i\}, \\ (m, j) & \text{if } \rho \text{ is obtained from } \pi \text{ by a merge step which combines} \\ & \text{components } A, B \text{ of } \pi \text{ and } j = \max(\min A, \min B). \end{cases}$$

Theorem 3.3. *The edge labelling λ for $\Pi_{\leq n}$ is an EL-labelling. Moreover, if $x < y$ in $\Pi_{\leq n}$ then the label sequences determined by λ for any two maximal chains in the interval $[x, y]$ have equal content.*

PROOF. First we show that the label sequence for every maximal chain C in $\Pi_{\leq n}$ has content $Q \setminus \{(m, 1)\}$. For $1 \leq i \leq n$ the label (c, i) appears in $\lambda(C)$ as the label for the covering relation $u \prec v$ which adds a component $\{i\}$. For $2 \leq j \leq n$, the label (m, j) appears as the label for the covering relation $u \prec v$ which merges two components in such a way that j is the minimal element of some component of u, but is not the minimal element of any component of v.

For any interval $[\pi, \rho]$ of height h in $\Pi_{\leq n}$, this implies there is a set of h distinct labels such that every maximal chain in the interval is assigned exactly these h labels. The only way for these labels to be ascending is for a maximal chain to first perform all necessary creation steps in ascending order, then at each subsequent step, merge the pair of components B_i, B_j which are both contained in the same block of ρ and which minimizes $\max(\min(B_i), \min(B_j))$. This also exhibits the existence of a rising chain. Since all label sequences have the same content, it follows that this rising chains has the lexicographically smallest label sequence among all maximal chains on the interval. $\square$

For any maximal chain $C \in \Pi_{\leq n}$ and any $i \geq 2$, (c, i) must appear before (m, i) in $\lambda(C)$, implying $\Pi_{\leq n}$ has no descending chains, hence yielding:

Corollary 3.4. *If $n > 1$ then $\Delta(\Pi_{\leq n} \setminus \{\hat{0}, \hat{1}\})$ is collapsible.*

There could perhaps be interesting S_n-module structure on the top homology of rank-selected subposets of $\Pi_{\leq n}$.

Corollary 3.5. *$\Pi_{\leq n}$ is supersolvable.*

PROOF. The EL-labelling of Theorem 3.3 easily translates to one in which the maximal chains are labelled by permutations in S_{2n-1}; simply replace each (c,i) by i and each (m,j) by $n+j-1$. This yields what McNamara calls an S_{2n-1}-EL-labelling in [Mc], implying $\Pi_{\le n}$ is supersolvable by a result in [Mc]. $\square$

The lattice $\Pi_{\le n}$ has also very recently been used in another context. It is shown in [BBH] to belong to a family of posets whose homology on lower intervals determine Poincare' series for free resolutions of a residue field k over rings $k[x_1,\dots,x_n]/I$ in which I is a monomial ideal; shellability of $\Pi_{\le n}$ together with the following Möbius function calculation are used in [BBH] to extend a Poincare' series denominator formula from [Be] to the degenerate case, namely the case in which I is the maximal ideal $(x_1,\dots,x_n)$ or more generally where some generators of I have degree one.

Corollary 3.6. *The Möbius function $\mu_{\Pi_{\le n}}(\sigma,\rho)=0$ unless each element in the support of ρ that is not in the support of σ forms a trivial component of ρ, in which case*

$$\mu_{\Pi_{\le n}}(\sigma,\rho)=(-1)^r\mu_{\Pi_m}(\sigma,\overline{\rho}),$$

where r is the number of trivial blocks in ρ that are not present in σ, Π_m is the partition lattice on elements in the common support of σ,ρ, and $\overline{\rho}$ is the restriction of ρ to this common support.

PROOF. It is straightforward to count, for the labelling λ given in the proof of Theorem 3.3, decreasing chains in each interval, to obtain the result. $\square$

4. Frames and cell-blocks in $PD_n(q)$

This section decomposes $PD_n(q)$ into subposets that are each isomorphic to $\Pi_{\le n}$.

Definition 4.1. *Call any set of n linearly independent lines, i.e.* 1-*dimensional subspaces, in F_q^n a* **frame**.

Each frame gives rise to a subposet of $PD_n(q)$ which we call a cell-block:

Definition 4.2. *Given a frame $l_1,\dots,l_n$, the corresponding* **cell-block** *is the subposet of $PD_n(q)$ consisting of all elements $\{W_1,\dots,W_r\}\in PD_n(q)$ such that each W_i is the span of some subset of $\{l_1,\dots,l_n\}$.*

Remark 4.3. *We will use the term cell-block both for this poset and for its order complex, since it will always be clear from context which is meant.*

Cell blocks are somewhat analogous to the notion of apartments in a building, though they are always balls rather than spheres. Every element of $PD_n(q)$ is contained in one or more cell-blocks. We define creation steps and merge steps in $PD_n(q)$ in a manner analogous to that in which they were defined for $\Pi_{\leq n}$.

Definition 4.4. *We say that $x = \{X_1, \ldots, X_k\}$ is obtained from $w = \{W_1, \ldots, W_{k-1}\}$ by a* **creation step** *if $dim(X_k) = 1$ and $X_j = W_j$ for $j < k$, and $z = \{Z_1, \ldots, Z_l\}$ is obtained from $y = \{Y_1, \ldots, Y_{l+1}\}$ by a* **merge step** *if $Z_l = Y_l \oplus Y_{l+1}$ and $Z_j = Y_j$ for $j < l$.*

Note that x covers w in $PD_n(q)$ if and only if x is obtained from w by a creation step or a merge step, and that each maximal chain in $PD_n(q)$ involves n creation steps and $n-1$ merge steps.

Remark 4.5. *The maximal chains in the cell-block determined by $\{l_1, \ldots, l_n\}$ are those whose creation steps create exactly the 1-spaces $\{l_1, \ldots, l_n\}$ in some order.*

Proposition 4.6. *Each cell-block in $PD_n(q)$ is isomorphic to the lattice $\Pi_{\leq n}$ of partial partitions.*

PROOF. Consider the cell-block determined by the frame $\{l_1, \ldots, l_n\}$. Map this to $\Pi_{\leq n}$ by sending each subspace $\langle l_{i_1}, \ldots, l_{i_k} \rangle$ to $\{i_1, \ldots, i_k\} \subseteq [n]$ to obtain the desired isomorphism. □

Our discrete Morse function on $PD_n(q)$ will build the order complex by sequentially attaching cell-blocks.

Definition 4.7. *Given a choice of an ordering on the cell-blocks, an* **open cell-block** *is the set of faces in a cell-block that are not shared with any earlier cell-block.*

5. Review: Notions from matroid theory

A few notions from matroid theory will provide a very convenient language for our upcoming analysis of the homotopy type of $PD_n(q)$. See [Ox] or [Bj2] for further background regarding matroids. The set of all lines in a finite vector space is the **ground set** of a matroid M. The **independent sets** of M are the sets of linearly independent lines. A **circuit** is a subset of the ground set which is not independent, with the property that deleting any element from it yields an independent set. A **basis** is a maximal independent set. Given a subset S of the ground set G, the **flat** generated by S, traditionally denoted $\overline{S}$, consists of those $g \in G$ such that either $g \in S$ or g together with some subset of S forms a circuit. We will usually denote by $\langle v_1, \ldots, v_n \rangle$ the flat generated by $\{v_1, \ldots, v_n\}$.

Let $<$ be any linear order on the ground set, in our case lexicographic order on vectors with respect to a fixed choice of coordinate system. Recall that a **broken circuit** is an independent set $\{v_{i_1}, \dots, v_{i_k}\}$ such that some $w \in \langle v_{i_1}, \dots, v_{i_k}\rangle$ satisfies $w < v_{i_j}$ for all $1 \le j \le k$. We will also use the following notion, which we have not seen in the literature.

Definition 5.1. A v_i**-broken circuit** is an independent set $\{v_1, \dots, v_k\}$ such that there exists $w \in \langle v_1, \dots, v_k\rangle$ with $w < v_i$ and $\langle v_1, \dots, v_k\rangle = \langle v_1, \dots, \hat{v}_i, \dots, v_k, w\rangle$ for some $1 \le i \le k$.

Each v_i which belongs to a v_i-broken circuit within a basis B is said to be **internally passive** in B because it may be exchanged for a smaller element to obtain an alternate basis.

Remark 5.2. *We will use v_i-broken circuits and internally passive elements to describe how cell-blocks overlap.*

If an element of a basis B is not internally passive, then it is **internally active**. The **internal activity** of B, denoted $ia(B)$ is the number of internally active elements of B.

6. Review: discrete Morse functions on poset order complexes and critical cell cancellation

See [Fo], [Ch], [He] for a review of discrete Morse theory in general and of lexicographic discrete Morse functions for poset order complexes. We follow terminology from [He]. We will use the following property of discrete Morse functions.

Remark 6.1. *If each interval in a graded poset admits a discrete Morse function on its order complex whose critical cells are all top-dimensional, then each interval is homotopy equivalent to a wedge of spheres of top dimension, implying the poset order complex is homotopically Cohen-Macaulay.*

Specifically, we will construct a discrete Morse function on each interval (x, y) in $PD_n(q)$ in such a way that all critical cells are top-dimensional, thereby proving Theorem 1.1. To this end, we use the lexicographic discrete Morse function of [BH]. See [He] for a succinct review of this material as well as some applications that are much easier than the one we will give shortly.

6.1. **Poset discrete Morse functions from non-lexicographic facet orders.** It is observed in [HeW, Section 3] that the lexicographic discrete Morse function construction of [BH] actually applies to any order complex facet order that satisfies the following condition, letting the rank of a poset element be its rank within the context of a chosen maximal chain:

Definition 6.2. Crossing condition. *Let $\le$ be a linear order on the maximal chains in a partially ordered set P of rank n with rank function* rk*. Let F be a maximal chain, $G \le F$ and $\sigma = F \cap G$. Suppose that $[n] - \{\text{rk}(p) \mid p \in \sigma\}$ is not an interval of natural numbers. Then there is some facet $G' \le F$ such that $F \cap G \subsetneq F \cap G'$.*

The crossing condition implies that for a maximal chain F, maximal faces in $F \cap (\cup_{G<F} G)$ are each supported on a set of ranks whose complement is a single interval of consecutive ranks. Thus it implies an interval system structure:

Definition 6.3. The set of faces in $F_j \setminus (\cup_{i<j} F_i)$ has an **interval system structure** if each maximal face in $F_j \cap (\cup_{i<j} F_i)$ omits exactly the ranks in one interval of consecutive ranks $[r+1, s-1]$. The interval system is the collection of these so-called I-intervals $[r+1, s-1]$ coming from the various maximal faces in $F_j \cap (\cup_{i<j} F_i)$.

Under these conditions, the faces in $F_j \setminus (\cup_{i<j} F_i)$ are exactly those faces of F_j that include at least one rank from every I-interval in the interval system of F_j, i.e. which avoid containment in any maximal face of $F_j \cap (\cup_{i<j} F_i)$. Obtain from the system of I-intervals a system of nonoverlapping J-intervals as follows:

Begin with a system of I-intervals, initializing the set of J-intervals to be empty. Repeat the following series of steps until there are no remaining I-intervals: (1) convert the lowest remaining I-interval, i.e. the one at lowest ranks, directly into a J-interval, (2) truncate the remaining I-intervals so that they do not share any ranks with any J-interval, and (3) discard any I-intervals that now strictly contain some other I-interval.

See [He] for a quick synopsis of how to obtain from an interval system structure a discrete Morse function such that

- each $F_j \setminus (\cup_{i<j} F_i)$ includes at most one critical cell
- $F_j \setminus (\cup_{i<j} F_i)$ includes a critical cell if and only if its system of J-intervals covers all ranks in F_j, in which case the dimension of the critical cell is one less than the number of J-intervals, with the critical cell consisting of those elements of F_j which appear as the lowest ranks within the J-intervals
- when the I-intervals already do not cover all ranks, then each face in $F_j \setminus (\cup_{i,J} F_i$ is matched by including/excluding the lowest rank not covered by any I-interval

See [He] for a description of the (slightly more complicated) matching when the I-intervals do cover all ranks.

Proposition 6.4 ([HeW]). *Let $F_1, \dots, F_k$ be a facet ordering on $\Delta(P)$ satisfying the crossing condition. Then each $F_j \setminus (\cup_{i<j} F_i)$ has an interval system structure, implying the acyclic matching construction of* [BH] *applies to $F_j \setminus (\cup_{i<j} F_i)$. Thus, $\Delta(P)$ has a discrete Morse function whose critical cells are described by interval systems as above.*

A **chain labelling** on a poset P is a labelling of its covering relations with elements of an ordered set Q, but with the label assigned to $u \prec v$ in a saturated chain $C = \hat{0} \prec u_1 \prec \cdots \prec u_k \prec u \prec v$ allowed to depend on the choice of saturated chain $u_1 \prec \cdots \prec u_k$ as well as depending on u, v.

Definition 6.5. Let P be a poset with the maximal chains ordered as follows. Given an edge labelling (resp. chain labelling) for P, group the saturated chains into classes based on label sequence content. Choose a total order on these content classes, then order maximal chains within each content class based on the lexicographic order on label sequences for maximal chains in the content class. The resulting order on facets in $\Delta(P)$ is a **content-lex facet order** if it satisfies the crossing condition, and for each interval (u, v) (resp. rooted interval $\hat{0} \prec u_1 \prec \cdots \prec u < v$) in P, the earliest maximal chain which includes both u and v (resp. which includes both $\hat{0} \prec u_1 \prec \cdots \prec u$ and v) is a rising chain when restricted to the interval (u, v).

When two critical cells in a discrete Morse function resulting from a content-lex facet order are in the same content class, then any gradient path from one to the other must stay within this content class at each step. In Section 7, we will provide a chain-labelling for $PD_n(q)$ giving rise to a content-lex facet order in which the content classes are the collections of maximal chains in the various cell-blocks.

6.2. **Gradient path reversal.** We will cancel pairs of critical cells within the same open cell-block. The fact that the labelling on a cell-block is an EL-labelling in which all maximal chains have equal content will allow us to construct a content-lex facet order in Section 7. To prove gradient path uniqueness, we will use the following result from [He], as reformulated for content-lex facet orders in [HeW].

Theorem 6.6 ([He]). *Let M be a generalized lexicographic discrete Morse function resulting from a content-lex facet order. Let $\tau^{(p+1)}$, $\sigma^{(p)}$ be critical cells coming from maximal chains in the same content class, and suppose that there is a gradient path γ from τ to σ. Suppose* *each downward step in γ prior to the last one swaps a pair of consecutive labels by deleting a chain element v such that the chain also has elements u, w with $u \prec v \prec w$. If the permutation transforming $\lambda(\tau)$ into $\lambda(\sigma)$ is* 321*-avoiding, then there is a unique gradient path from τ to σ.*

In order to cancel several pairs of critical cells simultaneously, the **multi-graph face poset**, denoted P^M, for the complex Δ^M of critical cells was defined as follows:

(1) The vertices in P^M are the cells in Δ^M, or equivalently the critical cells in the discrete Morse function on Δ.
(2) There is one edge between a pair of cells $\sigma^{(p)}, \tau^{(p+1)}$ of consecutive dimension for each gradient path from τ to σ.

Theorem 6.7 ([He]). *Any acyclic matching on P^M specifies a collection of gradient paths in $F(\Delta)$ that may simultaneously be reversed to obtain a discrete Morse function M' whose critical cells are the unmatched cells in the matching on P^M.*

7. A discrete Morse function for $PD_n(q)$ based on a content-lex facet order

In this section we construct a generalized lexicographic discrete Morse function on $PD_n(q)$ and describe its critical cells. They are not all top-dimensional, but in Section 8 we will cancel critical cells within each open cell-block to obtain a Morse function with only top-dimensional critical cells. Then in Section 9 we generalize this to arbitrary poset intervals to obtain the Cohen-Macaulay property. First let us describe the facet order that will give rise to our initial discrete Morse function.

Choose a fixed basis so as to specify a coordinate system. Also choose a fixed ordering on the elements of the ground field F_q such that $0 < 1 < x$ for all $x \neq 0, 1$, and use this to lexicographically order vectors expressed in our fixed coordinate system. Now identify the 1-spaces in F_q with their minimal elements, and order these minimal elements to obtain a total order on 1-spaces, i.e. on the ground set of our matroid. Now order frames as follows: list the 1-spaces $l_1, \ldots, l_n$ which generate a frame in lexicographically increasing order, then order these lists lexicographically to obtain a total order on the frames.

Next map any cell block to $\Pi_{\leq n}$ by $l_i \to i$. Finally, use the EL-labelling provided in Theorem 3.3 for $\Pi_{\leq n}$ to order maximal chains within each cell-block. This only gives a partial order on the set of maximal chains in a cell-block, because multiple maximal chains may receive the same label sequence. However, every linear extension

$F_1, \dots, F_k$ of this partial order on facets has the property that all maximal faces in each $F_j \cap (\cup_{i<j} F_i)$ belong to facets that come earlier than F_j in the partial order, because each maximal face in $F_j \cap (\cup_{i<j} F_i)$ results from a facet F_i which has a descent in F_j replaced by an ascent. From this we may conclude:

Remark 7.1. *All linear extensions of this EL-labelling on a cell-block will yield identical lexicographic discrete Morse functions for $PD_n(q)$, so there is no need to specify a particular linear extension for each cell-block.*

Thus, we may order maximal chains by first ordering cell-blocks by the above order on frames, then within each cell-block using the EL-labelling for $\Pi_{\le n}$ from Theorem 3.3 to order its maximal chains lexicographically. Proposition 7.4 will show that this facet order satisfies the crossing condition, so the [HeW] generalization of the [BH] lexicographic discrete Morse function construction will apply.

Definition 7.2. The restriction of a maximal chain M to an interval (u, v) is **delinquent** if M has ascending labels on that interval, but there is another maximal chain that agrees with M except strictly between u and v and which belongs to an earlier cell-block. It is **minimally delinquent** if there is no $u' < v'$ in M with $u \le u' < v' \le v$ and $(u', v') \ne (u, v)$ such that M is also delinquent on (u', v').

The following will be useful for analyzing overlap between cell-blocks.

Proposition 7.3. *If a maximal chain M is minimally delinquent on the interval (u, v), then it begins with a covering relation $u \prec u'$ which creates a 1-space l_i. Moreover, l_i must be internally passive within some component of v, with respect to the restriction of the frame of M to this component of v; in addition, v must be the lowest element of M in which l_i is internally passive.*

PROOF. For a chain to be shared with an earlier frame, one or more of the basis elements for the later frame must be exchangeable for earlier elements of the ground set. In particular, the jump from u to v, must enlarge some component of u by adding vectors not in the collective span of the components of u, enabling distinct saturated chains from u to v to create distinct 1-spaces and still all lead to v. Each saturated chain must create the same number of 1-spaces to be added to each component of v and the same number of 1-spaces for new components. Furthermore, creation steps must precede merge steps in M for its labels to be ascending. There will be a delinquent chain from u' to v for u' the highest element of M such that the covering relation upward from M is a creation step that is internally passive in some component of v with respect to the basis for the frame of M restricted to this component. Thus, the result follows from minimality of the delinquent chain. □

Proposition 7.4. *The above facet order on $PD_n(q)$ satisfies the crossing condition.*

PROOF. Let us show for each F_j that any $\sigma \in F_j \cap (\cup_{i<j} F_i)$ which skips multiple rank intervals is strictly contained in some $\tau \in F_j \cap (\cup_{i<j} F_i)$. If F_j has a descent at rank r, then there exists some $i' < j$ such that $F_{i'}$ lies in the same cell block as F_j with $F_{i'}$ and F_j differing only at rank r and $\lambda(F_{i'})$ having the same content as $\lambda(F_j)$.

Thus, we may assume σ skips only ranks where F_j is increasing. Now consider the lowest delinquent chain (u, v) in F_j skipped by σ. We may assume it is a minimally delinquent chain, since otherwise the desired τ may be obtained by skipping fewer ranks on this interval. Thus, $[u, v] = u = u_0 < u_1 < \cdots < u_r < u_{r+1} < \cdots < u_k = v$ where for $1 \le i \le r$ the element u_i is obtained by creating l_i with $l_1 < \cdots < l_r$ and for $r+1 \le i \le k$ the element u_i is obtained from u_{i-1} by a merge step. Moreover, l_1 is internally passive in the component of v that contains it, but not internally passive in the component of v_{k-1} that contains it. This implies there must be an earlier facet F_i also containing σ, which instead has creation steps $v'_1, \ldots, v'_r$ with $v'_1 < \cdots < v'_r$, followed by its own series of merge steps. Let B be the basis determining the cell block containing F_j. Then $B' = B - \{v_1, \ldots, v_r\} + \{v'_1, \ldots, v'_r\}$ is a lexicographically smaller basis, since v_1 is the smallest vector in $B \setminus B'$, but $v'_1 \in B' \setminus B$ is smaller than v_1. Let $F_{i'}$ agree with F_j except on the interval skipped by σ, where $F_{i'}$ instead coincides with F_i. $F_{i'}$ has basis B', so precedes F_j in our facet order. Let $\tau = F_j \cap F_{i'}$, so we are done. □

Proposition 7.5. *The above facet order is a content-lex facet order.*

PROOF. This is immediate from the following chain labelling, together with Proposition 7.4 and the fact that its restriction to any particular cell-block is equivalent to the EL-labelling of Theorem 3.3. Label each covering relation $x \prec y$ which creates a new 1-dimensional space, by the pair (c, v) where v is the vector in this 1-dimensional space which has a 1 as its first nonzero coordinate. Let $l_1, \ldots, l_i$ be the list of 1-spaces which were created in the saturated chain from $\hat{0}$ to x; then label each covering relation $x \prec y$ which replaces a pair of vector spaces

V_1, V_2 by their direct sum with the label (m, v) where v_j is the lexicographically smallest vector in $\{\langle l_1\rangle, \langle l_2\rangle, \cdots, \langle l_i\rangle\} \cap V_j$ for $j = 1, 2$ and $v = \max(v_1, v_2)$. Now order labels by $(m, v) > (c, w)$ for all v, w, and use the linear order on the second coordinate when two labels agree in the first coordinate. It is easy to check that ordering content classes in this labelling as described above then using any linear extension of lexicographic order within each content class gives the facet order already described in this section. $\square$

Having a content-lex facet order will allow us to use Theorem 6.6 to cancel pairs of critical cells within the same open cell-block. Paired critical cells will have the same content.

As an example of a critical cell, consider the ground set $v_1 = (0, 1), v_2 = (1, 0), v_3 = (1, 1)$ for $PD_2(2)$ with the ordering $v_1 < v_2 < v_3$. Then the saturated chain labelled $(c, v_3), (c, v_2), (m, v_3)$ has a descent followed by a minimal delinquent chain, so it contributes a critical cell of dimension one. The delinquent chain results from the v_2-broken circuit $\{v_2, v_3\}$ which allows the earlier increasing chain labelled $(c, v_1), (m, v_3)$ on the interval from the 1-space l_3 to the full 2-dimensional space.

8. Matching and cancelling pairs of critical cells

In this section, we cancel all but some top-dimensional critical cells in the Morse function of the previous section, to show $PD_n(q)$ has the homotopy type of a wedge of spheres of top dimension. We always cancel pairs of critical cells contributed by maximal chains from the same cell-block. Proposition 7.5 allows us to use Theorem 6.6 to prove gradient path uniqueness, as needed to cancel pairs of critical cells. The descriptions in this section of these gradient paths to be reversed are rather terse, but readers are referred to [He] and [HeW] for much more detailed examples of gradient paths that shift labels in a very similar manner.

The approach will be to consider an arbitrary cell-block, and to construct an acyclic matching on the restriction of P^M to this open cell-block in such a way that only top-dimensional critical cells from the open cell-block are left unmatched. The Cluster Lemma (cf. [Jo, Lemma 2]) as extended slightly in [He] implies that the union of these acyclic matchings on various portions of P^M is acyclic on P^M. Theorem 6.7 shows that any acyclic matching on P^M allows us the simultaneous reversal of the gradient paths given by the matching, and hence the cancellation of all the matched critical cells. The end result will be a discrete Morse function on $PD_n(q)$ with only top-dimensional critical cells, as desired. Thus, the main focus of this section is to show how to match critical cells within any one open cell-block.

Let $B = \{l_1, \ldots, l_n\}$ be the basis of linearly independent lines giving rise to the cell-block. Let $v_1, \ldots, v_n$ be the vectors in the lines $l_1, \ldots, l_n$ which have 1's as their leading nonzero coordinates, with vectors listed in lexicographically increasing order. We will speak of l_i and v_i interchangeably.

In an effort to simplify notation in the (fairly complex) upcoming argument, we will sometimes use c_i to denote a creation step, namely the covering relation in which the new one-dimensional space l_i is created. We use c_i, c_j (resp. m_i, m_j) with $i < j$ denotes a pair of creation steps (resp. merge steps), where the former has smaller label than the latter in the chain-labelling of Proposition 7.5.

8.1. **Cell-blocks with nonzero internal activity.** This section deals with a class of open cell-blocks for which there is a fairly simple complete acyclic matching, implying that attaching these cell-blocks does not change the homotopy type of the complex. The argument below is subsumed by the general case, but is included both as a "warm-up" for the more intricate proof needed for arbitrary cell-blocks, and to demonstrate the strong analogy with the standard shelling for matroid complexes.

Theorem 8.1. *If B has nonzero internal activity, then all critical cells in the open cell-block given by B may be cancelled.*

PROOF. Let v_i be the smallest internally active element of B. Notice that there must be a merge step above $c_i := (c, v_i)$, since the 1-space spanned by v_i must itself be merged at some point. This ensures an ascent above c_i in every maximal chain. Therefore, every maximal chain contributing a critical cell must have at least one delinquent chain above c_i. Notice that c_i can never appear as the bottom of a delinquent chain, since it is internally active in B, so it appears either in the interior of a delinquent chain or within a descending series of labels.

Match any critical cell τ contributed by a maximal chain M that has c_i as part of a descending series S of consecutive labels just below a delinquent chain D by shifting c_i upward to its unique ascending position within the label sequence on D. This will match all critical cells, since each has c_i either in a descending series S or in a delinquent chain D just above such a (possibly empty) descending series. It is not hard to construct a gradient path that shifts c_i in this fashion, ensuring this indeed gives a matching on P^M. Simply let each successive downward step in the gradient path delete the chain element immediately above c_i, allowing c_i to shift upward past a label that is smaller than it, then the subsequent upward gradient path step will insert a chain element just below c_i. The result is a maximal chain M'. Next we will show that M' contributes a critical cell σ of dimension $dim(\tau) - 1$.

If c_0 is the creation step initiating D, and m is the merge step concluding D, then $c_0 < c_i < m$. Inserting c_i into D does not impact the minimality of the delinquent chain from c_0 through m, since c_i is internally active. Note that M' has one fewer descent than M and has the same number of minimally delinquent chains, so M' has one fewer minimal skipped interval than M. Furthermore, the truncation procedure on minimal skipped intervals will not impact this relationship in dimensions. Theorem 6.6 implies that the gradient path described above is indeed a unique gradient path from τ to σ. Thus, we obtain a complete matching on the restriction of P^M to this open cell-block.

Acyclicity follows from the fact that P^M restricted to our open cell-block has a product structure $P^M = Q \times C_2$ where C_2 is a chain with 2 elements where we match $(x, \hat{0})$ with $(x, \hat{1})$ for any $x \in Q$. By Theorem 6.7, this complete acyclic matching on P^M enables cancellation of all critical cells contributed by the cell-block, by simultaneously reversing the gradient paths represented by matching edges. □

Recall that in the standard shelling for a matroid independence complex, only the bases with internal activity 0 change the homotopy type of the complex with their attachment (cf. [Bj2]). We have established the following analogous situation.

Corollary 8.2. *Let $A_1, \dots, A_s$ be the ordering on cell-blocks resulting from lexicographically ordering frames. Then the homotopy type of the partial complex $A_1 \cup \dots \cup A_r$ for $r < s$ does not change with the attachment of cell-block A_{r+1} if the basis specifying A_{r+1} has nonzero internal activity.*

Now we turn to cell-blocks with internal activity zero, i.e., the much harder case.

8.2. **Terminology, notation and general strategy for matching.** The remainder of Section 8 is devoted to matching and cancelling all non-top-dimensional critical cells in any fixed open cell-block which has internal activity zero. We begin with notation and an outline of the general matching strategy to be used. For convenience, we will often speak of a maximal chain and its label sequence (when this determines its series of covering relations) interchangeably.

In light of our description of which maximal chains contribute critical cells, the following definition will be essential to our matching.

Definition 8.3. A **λ-block** in a maximal chain is a segment of the maximal chain which is labelled by a maximal sequence of consecutive labels of the form

$$m_{r_1} > \cdots > m_1 > c_{r_2} > \cdots > c_1 > c_0^I < \cdots < c_{r_3}^I < m_1^I < \cdots < m_{r_4}^I,$$

excluding m_{r_1} if it may belong to a lower λ-block so as to ensure that λ-blocks are disjoint. The highest λ-block in a maximal chain might consist only of $m_{r_2} > \cdots > m_1$ while every other one must include labels c_0^I and $m_{r_4}^I$.

The label sequence on any maximal chain contributing a critical cell consists of a series of consecutive λ-blocks. We will often speak about matching a λ-block B by modifying it in a particular way, by which we actually mean matching a critical cell contributed by a maximal chain containing B with the critical cell from another maximal chain whose label sequence is obtained by modifying B in the specified way and otherwise leaving the label sequence unchanged.

Notation 8.4. *Given a creation step c_i which is internally passive in a maximal chain M, denote by $m(c_i)$ the earliest merge step $u \prec v$ in M which causes c_i to be internally passive in some component of v, i.e. which "reduces" c_i. If a merge step reduces both c_i and c_j, denote this by $m(c_i, c_j)$.*

By Proposition 7.3, each minimally delinquent chain begins with a creation step c and concludes with $m(c)$. We say that c **initiates** this delinquent chain.

Notation 8.5. *If a λ-block B has a single minimally delinquent chain, denote this by $del(B)$. When B potentially has more than one overlapping minimally delinquent chain, denote this $dels(B)$.*

Lemma 8.6. *A maximal chain contributes a critical cell if and only if it is covered by λ-blocks with $dels(B)$ containing at most two delinquent chains for each λ-block B.*

PROOF. For the I-intervals to cover a maximal chain, the chain must consist of a series of λ-blocks, so as to be covered by descents and delinquent chains. Notice for any $dels(B)$ that its labels will be ascending with each minimally delinquent chain beginning with a creation step and ending with a merge step. Therefore the lowest elements of all minimally delinquent chains in $dels(B)$ are below the highest elements of all minimally delinquent chains in $dels(B)$, implying that there is at least one chain element covered by all of them. The truncation procedure for converting I-intervals to J-intervals will include the lowest of these minimally delinquent chains in its entirety, the truncated portion of the next lowest one and will eliminate all of the others, because they violate minimality after truncation. This leaves some ranks uncovered unless there are at most two minimally delinquent chains in $dels(B)$.

Notation 8.7. *Denote the delinquent chain initiated by c as $del(c)$.*

Definition 8.8. *In an effort to increase readability, we introduce some descriptive but rather informal sounding terminology: we say that v_i is* **needed to reduce** *v_j in a basis I if $v_i, v_j \in I$, v_j is internally passive in I, but v_j is internally active in $I \setminus \{v_i\}$.*

We adopt this somewhat informal sounding terminology in the hopes that it will help increase the readability of the very technical upcoming proof. Given a sequence $m_{r_1}, \dots, m_1, c_{r_2}, \dots, c_1$ of labels with descending values in a λ-block B, then we say that the *ascending position* of a label μ within $dels(B)$ is the unique place where μ could be inserted without introducing any descents, if such a position exists. Likewise, the *descending position* for μ within a series of descending labels is the unique position it could be inserted without introducing any ascents.

Definition 8.9. *A creation or merge step is said to be* **unanchored** *if it is neither the bottom nor the top label of any minimal delinquent chain. An unanchored creation or merge step is* **shiftable** *if it may be shifted via a gradient path from its ascending position within a delinquent chain $dels(B)$ to a descending position in B just below $dels(B)$ to obtain a critical cell, or it may be shifted from a descending position just below $dels(B)$ to an ascending position within $dels(B)$ to obtain a critical cell.*

Definition 8.10. *A creation step (resp. merge step) within a delinquent chain is an* **unanchored creation step** *(resp.* **unanchored merge step***) if it does not initiate or conclude that delinquent chain.*

To eliminate all non-top-dimensional critical cells means exactly that the minimal delinquent chains in critical cells surviving cancellation must each consist of a covering relation c_i immediately followed by $m(c_i)$. This section will match and cancel pairs of critical cells by gradient path reversal, using four rounds of matching described informally as follows, with later rounds applying to cells not matched in earlier rounds.

- round one: match by shifting the smallest overall shiftable creation step from a descending position just below some $dels(B)$ to its ascending position within $dels(B)$ or vice versa, thereby eliminating all critical cells in which any minimal delinquent chain has any internal creation steps.
- round two: match based on the lowest possible λ-block B which is matchable either via the Splitting Map (see Section 8.5), or by shifting an unanchored merge step m from its descending position just below $dels(B)$ to its ascending position within $dels(B)$ or vice versa.
- round three: match all λ-blocks B not yet matched which have a single minimal delinquent chain with one or more internal merge steps
- round four: match the following two remaining types of critical cells: (1) ones with some $dels(B)$ having two overlapping delinquent chains $del(c_1), del(c_2)$ with $c_1 < c_2$ and with some c below $dels(B)$ satisfying $m(c) \in del(c_2)$ and $m(c) > m(c_1)$, and (2) ones with consecutive λ-blocks B', B where the Splitting Map S_2 applies to B yielding a critical cell with consecutive λ-blocks B', B_1, B_2 which is instead matched by applying S_2^{-1} to (B', B_1).

The remainder of this section describes and justifies in detail these four rounds of critical cell cancellation. Round one handles all critical cells having any creation steps in the interior of minimally delinquent chains; the later rounds deal with merge steps in the interior of delinquent chains and with overlapping delinquent chains. Together these rounds will cancel all non-top-dimensional critical cells.

Results from [He], as reviewed in Section 6, will be used to show that there is a unique gradient path between each pair of critical cells to be matched and cancelled. Acyclicity of the matching for each round is checked, working in terms of the multigraph face poset P^M consisting of the critical cells remaining at that stage, with suitably adjusted covering relations. Later rounds sometimes involve reversing matching edges from earlier rounds. The end result will be an acyclic matching with only top-dimensional critical cells for any particular open cell-block. Since our original discrete Morse function comes from a filtration, the acyclic matchings on the various open cell-blocks combine to yield a global acylic matching.

8.3. **Two key lemmas.** The first lemma below seems to be at the heart of why $PD_n(q)$ is Cohen-Macaulay. Section 12 will generalize this to a property possessed by certain matroids, called Property R, to give some suggestion of how our results might perhaps generalize to partial decomposition posets for other matroids.

Definition 8.11. *We say that $v \in I$ is* **internally passive** *with respect to an independent set I if v may be replaced by a smaller element to obtain a new independent set with the same span as I.*

Lemma 8.12 (Lemma R)**.** *Let $\langle v_1\rangle, \langle v_0\rangle$ be 1-spaces in a basis B within which both are internally passive. Then at least one of the following must hold:*

(1) *v_1, v_0 have leading ones in the same position*
(2) *v_1 is not needed to reduce v_0*
(3) *v_0 is not needed to reduce v_1*

PROOF. Suppose v_1 has leading one strictly farther to the left than v_0 and is needed in the reduction of v_0. Then v_1 itself must be reduced by vectors in B other than v_0 before being capable of reducing v_0. □

Corollary 8.13. *Any internal creation steps present within a minimal delinquent chain initiated by c_0 are not needed to reduce c_0.*

Lemma 8.14 (Lemma L)**.** *Given overlapping minimal delinquent chains initiated by c_0, c_1, respectively, then v_1 is not needed to reduce v_0.*

PROOF. If v_0, v_1 have leading ones in the same position, then $\langle v_0\rangle, \langle v_1\rangle$ cannot be merged prior to $m(c_0)$; otherwise, the delinquent chain $del(c_0)$ would not be minimal. If v_1 has larger leading one than v_0, then v_1 cannot be used in the reduction of v_0 without itself first being reduced, but $m(c_1)$ must come strictly after $m(c_0)$. Thus, v_1 again is not needed to reduce v_0. □

8.4. **Critical cell matching: round one.** Suppose there is an unanchored creation step c in a λ-block B which is shiftabe from below $dels(B)$ to within $dels(B)$ or vice versa; that is, suppose there is a gradient path accomplishing this. Then c must be intermediate in value to the smallest and largest labels of $dels(B)$; furthermore, $dels(B)$ must not contain $m(c)$ unless $m(c)$ concludes $dels(B)$ and $m(c) = m(c')$ for some $c' > c$ appearing within $dels(B)$. Notice that Lemma 8.12 implies that any creation step appearing in the interior of $dels(B)$ either itself initiates a minimally delinquent chain or else meets these conditions.

Then match the critical cell based on the smallest such c, i.e., match by choosing the smallest such c, then match by shifting c from just below some $dels(B)$ to its ascending position within $dels(B)$, or vice versa. By Lemma 8.12, this matches all critical cells with any internal creation steps. Gradient path existence and uniqueness follow exactly as in the previous section.

Proposition 8.15. *This first round of matching on P^M is acyclic.*

PROOF. All matching edges in P^M preserve the set of minimal and maximal elements of the delinquent chains and shift creation steps downward from within delinquent chains. Suppose there were a cycle C. Consider a matching step in C which shifts the smallest c which is ever moved within C to just below some $dels(B)$. This is followed by a downward step that leaves c below $dels(B)$, and now we are at the top of a matching edge, precluding continuation of the cycle. ☐

Now what remains is to deal with internal merge steps and overlapping delinquent chains.

8.5. **The Splitting Map on λ-blocks.** In preparation for the next round of cancellation, we introduce a "splitting map", denoted S, which takes as input a λ-block B having either of the two forms described shortly, and outputs a pair of λ-blocks (B_1, B_2). Denote S by S_1 or S_2, depending on which form B takes. This definition only applies to B appearing in maximal chains contributing critical cells that were not already matched in round one.

Definition of S_1: Suppose B has a unique minimal delinquent chain, letting c_0 initiate $del(B)$. Let $c_0, \dots, c_m$ be the creation steps of B. Let c_i be the smallest unanchored creation step below $del(B)$ in B with the following properties, if one exists:

(1) $m(c_i) \in del(B)$
(2) c_i can be reduced in V without using any of the smaller creation steps, namely $c_0, c_1, \dots, c_{i-1}$, where V is the subspace of $\mathbb{F}^n$ containing v_i at the conclusion of B

Then S_1 sends B to a pair of consecutive λ-blocks (B_1, B_2) with c_i initiating $del(B_1)$ and c_0 initiating $del(B_2)$. Let j be the smallest index such that c_i can be reduced within B without using any c_k with $k > j$ or $k < i$. The labels in $del(B_1) \cup del(B_2)$ are exactly the set of labels in $del(B)$ along with c_i. Then B_1 consists of c_i along with:

(1) those creation steps c such that (a) $c_i \leq c \leq c_j$ and (b) c can be reduced in B without using any c_k with $k < i$ or $k > j$
(2) the lexicographically smallest choice (when listed in ascending order) of merge steps from B needed to reduce all 1-spaces created in B_1

The following key properties are not hard to verify.

Proposition 8.16. *Suppose S_1 yields (B_1, B_2). Then all creation steps included in B_1 are reducible within B_1, and all merge steps m included in B_1 satisfy $m \le m(c_i)$.*

Definition of S_2: Suppose B has exactly two overlapping minimally delinquent chains, initiated by c_0, c_1, respectively. Moreover, suppose each c appearing below $dels(B)$ in B either has $m(c) \in del(c_0)$ or else satisfies $c > c_1$ and $m(c) = m(c_1)$. Then S_2 splits B into (B_1, B_2) with delinquent chains initiated by c_0, c_1, respectively, and the set of labels in $dels(B_1) \cup dels(B_2)$ equalling exactly the set of labels in $dels(B)$. Included in B_1 are exactly the same choice of creation steps and merge steps as was made for S_1, with the exception of c_0, c_1 as already specified. That is, apply the above rules now, treating c_0 now as c_1 is treated in S_1. Lemma L ensures that such splitting is possible, and that S_2 yields a maximal chain contributing a critical cell.

Remark 8.17. *When B has exactly two overlapping minimal delinquent chains, then any creation step $c \in B$ below $dels(B)$ which fails the requirements of S_2 above will cause the critical cell to get matched either in round one or in round four.*

The definitions of S_1 and S_2 easily yield the next two propositions:

Proposition 8.18. *If a critical cell is not matched in round one, then applying S yields a critical cell also not matched in round one.*

This follows from the fact that $c \in B_1$ implies $m(c) \in del(B_1)$.

Proposition 8.19. *S is injective as a map on λ-blocks.*

Remark 8.20. *S is not, however, injective on label sequences for entire maximal chains; there may be two distinct maximal chains where S applies to entirely different λ-blocks in these two maximal chains to yield the same image maximal chain.*

Lemma 8.34 will show that critical cells not matched by applying S because of this fact will instead be matched by virtue of having a pair of consecutive λ-blocks in $im(S)$, except when unsplitting causes there to be three overlapping minimal delinquent chains; in that case, the critical cells will be matched in round four. Lemma 8.34 will follow easily from choices made in the definition of the splitting map.

Denote by S^{-1} the map that merges the lowest pair of consecutive λ-blocks $(B_1, B_2) \in im(S)$ in a maximal chain M as follows, letting B denote $S^{-1}(B_1, B_2)$. Obtain $dels(B)$ by arranging all creation steps and merge steps that occur in $dels(B_1) \cup dels(B_2)$ into the lexicographically smallest achievable ascending label sequence, excluding from $dels(B)$ the label initiating $dels(B_1)$ if it is larger than the label initiating $dels(B_2)$. We sometimes denote $S^{-1}(B_1, B_2)$ by $B_1 \cup B_2$.

8.6. **Critical cell matching: round two.** Next match each remaining critical cell based on its lowest λ-block B that is matchable in either of the following two ways, when such a B is present:

(1) B has a collection $B_1, \dots, B_{i+1}$ of consecutive λ-blocks immediately above B, listed from lowest to highest location in the chain, for i as small as possible such that $(B, B_1, \dots, B_{i+1}) \notin im(S^{i+1})$. If i is even and S is applicable to B, then B is matchable by applying S to B. If i is odd for $i > 1$, then B is matchable by applying S^{-1} to (B, B_1).
(2) B has internal merge steps, letting m be the smallest merge step achievable in B by merging blocks not created in B which get merged in B. Then B is matchable by shifting m from the interior of $dels(B)$ to just below $dels(B)$ or vice versa, unless m is larger than either the largest label in the delinquent chain just below B or the maximal label in $dels(B)$, follwing the conventions of Remark 8.23 .

Notation 8.21. *Denote by M the map sending a critical cell to its matching partner.*

Match via option (1) if possible for the lowest B of either type; otherwise match by option (2) for the same B. In option (1), if $i = 0$ it may happen that $S(B)$ is already matched while B is not, due to the lower λ-block in $S(B)$ being in the image of the splitting map, together with the λ-block below it, but then B will be matched in round four.

We now check that matching option (2) is well-defined by showing that the map shifting m downward to its descending position below $dels(B)$ is injective and then making a choice for how to shift downward.

Lemma 8.22. *Shifting a merge step labelled m downward past a merge step labelled m' for $m' < m$ is an injective map.*

PROOF. Let $m' = \max(\min(V_1 \cap I), \min(V_2 \cap I))$, for V_1, V_2 the vector spaces being merged and I our frame. Let W be the vector space obtained by merging V_1 and V_2. There is a unique choice how to swap m, m' unless m merges W with another vector space V_3. But $m > m'$ implies $\min(V_3 \cap I) > \max(\min(V_1 \cap I), \min(V_2 \cap I))$, and there are two ways to swap the order of m, m', as follows: when m appears below m', it may either merge V_1, V_3 or merge V_2, V_3, and we get an injection. $\square$

Remark 8.23. *When applying Lemma 8.22 to shift m downward past m', we make the convention that m merges V_2, V_3, for $\min(V_1 \cap I) < \min(V_2 \cap I)$, though this choice is not important. Since merge steps appearing between delinquent chains do not necessitate matching, it is not necessary (and indeed is not true) that shifting m upward past m' into $dels(B)$ also gives an injection.*

Proposition 8.24. *In the second round of matching on critical cells, $M(C) = C'$ if and only if $M(C') = C$, except when S_2 applies to B to yield (B_1, B_2) such that B_1 together with the block immediately below it are in $im(S_2)$.*

PROOF. Applying S or S^{-1} to C cannot cause additional unanchored creation steps to become shiftable, and it can only cause a creation step c' to switch from initiating a delinquent chain to being unanchored as follows: $dels(c')$ must be merged with the delinquent chain above it by S_1^{-1}, in which case $m(c')$ appears in the merged delinquent chain, so c' still is not shiftable. Whether S^{-1} takes the form S_1^{-1} or S_2^{-1} is determined by whether the creation steps initiating consecutive delinquent chains appear in decreasing or ascreasing order. Whether to apply S or S^{-1} to B is determined by the parity of i, for $B \in im(S^i) \setminus im(S^{i+1})$. Thus, if $M(C)$ is obtained by applying S, the only way to obtain C' which matches with a cell other than C is when S is not injective, so that $S^{-1}(C') \neq C$.

If we match by option (2), then Lemma 8.22 specifies how to do this. To see that the partner cell under the injection is not matchable in a preferable way, note that this shifting of a merge step will not impact which λ-blocks are splittable, belong to $im(S)$, or have shiftable creation steps. $\square$

Remark 8.25. *The excluded critical cells in Proposition 8.24 will be matched in round four.*

Proposition 8.26. *If a critical cell C is matched by applying S_1 to a λ-block B, then the resulting maximal chain also contributes a critical cell, denoted D.*

PROOF. Let $S_1(B) = (B_1, B_2)$, with B_1 and B_2 having delinquent chains labelled $c_1, \dots, m(c_1)$ and $c_2, \dots, m(c_2)$, respectively. We have $c_1 > c_2$, and must check that $m(c_1), m(c_2)$ are the largest merge steps in B_1, B_2, respectively. For $m(c_1)$, this follows from the necessity of all merges included in B_1. For $m(c_2)$, this follows from the fact that the largest label in $dels(B)$ was the first to reduce c_2 there. $\square$

The next proof is given in quite a bit of detail, providing an example of how these types of arguments; similar arguments later will be left to the reader, as they are completely straight-forward.

Proposition 8.27. *The critical cells in Proposition 8.26 satisfy* $\dim(D) = \dim(C) + 1$*, and there is a unique gradient path from* D *to* C*.*

PROOF. It is easy to see that applying S_1 to B yields a maximal chain with one more J-interval than the original maximal chain had, with J-intervals still covering all ranks. Let us now construct very explicitly a gradient path γ from D to C. Let c_1, c_2 initiate $dels(B_1), dels(B_2)$, respectively. First consider the case where $m(c_1)$ immediately precedes c_2. Take a downward step by deleting the chain element separating $m(c_1)$ from c_2, causing c_2 to shift downward to below all labels except c_1 in $del(c_1)$. Then take the matching step inserting a chain element just above c_2, since there is an ascent at this rank which is not in any I-interval, while all lower ranks do belong to I-intervals. Next take a downward step deleting the chain element just below $m(c_2)$, yielding C. Uniqueness of this gradient path follows from the argument used to prove Theorem 6.6.

Now suppose there are creation steps situated between $m(c_1)$ and c_2. If any such c_i is larger than c_1, then applying S_1^{-1} to the pair of λ-blocks would need to shift c_i below c_1, introducing an inversion between c_1 and c_i. This can only be accomplished by traversing a matching edge from the first round of critical cell cancellation in an upward direction. Thus, cancelling a pair of critical cells via the Splitting Map involves reversing some face poset matching edges from the previous round. The gradient path γ must proceed sequentially from largest to smallest such c_i, shifting each in turn to below $del(c_1)$; to see this, note that after some c_i has been shifted to below c_1 and some $c_{i'} < c_i$ is shifted downward into $dels(c_1)$, the matching would shift $c_{i'}$ to below $dels(c_1)$ rather than shift c_i back into $dels(c_1)$, since $c_{i'} < c_i$. For each individual c_i, γ first swaps c_i with the merge step $m(c_1)$ immediately below it, then has an upward (matching) step shifting c_i to immediately below c_1. After all this downward shifting of creation steps, γ has a downward step combining $del(c_1)$ and $del(c_2)$ into a single delinquent chain initiated by c_2. Gradient path uniqueness follows from the fact that the label sequence content must be preserved, together with the fact that downward steps cannot introduce inversions. □

Propositions 8.26 and 8.27 yield the following, in the case where the matching M results from applying S_1. Other cases are similar, and left to the reader.

Proposition 8.28. *If C is critical of dimension d, then $M(C)$ is critical of dimension $d \pm 1$.*

To avoid confusion in the following argument, we refer to covering relations in the multi-graph face poset as **transitions** and covering relations in $PD_n(q)$ as **steps**.

Proposition 8.29. *This second round of matching is acyclic.*

PROOF. We need only consider $P^{M'}$ for M' the set of critical cells not matched in round one. Suppose there is a cycle C. Then$C$ must alternate upward and downward transitions. Suppose first that C has no matching transitions resulting from applying the splitting map. Then each upward-oriented transition shifts a merge step downward from the interior of a delinquent chain to below it, and all transitions preserve the set of delinquent chains. Consider the upward transition shifting a merge step to a lower position than is done by any other transition in C. This is followed by a downward transition which must place us at the top of a matching edge, precluding continuation of the cycle.

If C includes transitions which apply the splitting map, consider the one applying to the lowest λ-block. This is followed by a downward transition putting us at the top of an upward-oriented transition, by virtue of the block just split, unless the block splitting enables application of the splitting map or the shift of a merge step in a still lower λ-block. But the way the splitting map was defined makes it impossible that the lower split block or a still lower one has an unanchored merge step that will shift upward into a delinquent chain from below if the unsplit block and those below it did not also have this. We also cannot change splitting status of lower blocks, so the downward transition must leave us at the top of an upward-oriented edge, unable to complete the cycle. $\square$

8.7. **Restrictions on cells not yet matched.** Next we establish rather severe restrictions on the critical cells not yet matched.

Proposition 8.30. *Suppose a critical cell not yet matched has consecutive λ-blocks (B', B) with internal merge steps in $dels(B)$, all of which are larger than the label concluding $dels(B')$ and has no creation steps below $dels(B)$ in B., Also suppose no lower λ-blocks have merge steps enabling matching by shifting them. Then $(B', B) \in im(S)$, or S applies to B.*

PROOF. The fact that internal merge steps are too large for downward shifting exactly enables the two λ-blocks to be combined into a single one. This amounts to applying S^{-1}, unless there is at least one creation step between B' and B that would instead appear in B' when the splitting map is applied to the merged block. But one may easily check that any such creation step would be of a form that makes the splitting map applicable to B. Using the fact that the critical cell was not matched based on any lower λ-block, our greedy choices in the definition of the splitting map are exactly as needed here. □

Next consider an internal merge step m in some $dels(B)$ which is not shiftable to below $dels(B)$ in the λ-block because there is some creation step c below $dels(B)$ such that m merges two vector spaces, one of which contains the 1-space created by c. Let B be the lowest λ-block having an internal merge step. Then B has label sequence $c_{j_r}, c_{j_{r-1}}, \ldots, c_{j_1}, c_i, m_1, \ldots, m_s, m(c_i)$ for some $r \geq 1, s \geq 1$, with some label among $c_i, c_{j_1}, c_{j_2}, \ldots, c_{j_{i'-1}}$ needed in the reduction of each $c_{j_{i'}}$, and with $m(c_{j_{i'}}) \in dels(B)$ for each $j_{i'}$. In the arguments below, v_j denotes an element of $\{v_{j_1}, \ldots, v_{j_r}\}$. Also make the convention that a pair of critical cells to be matched are denoted by σ, τ when we have $\sigma \subseteq \overline{\tau}$ (as opposed to having $\tau \subseteq \overline{\sigma}$).

Proposition 8.31. *Suppose a critical cell not yet matched has internal merge steps and also one or more creation steps below $dels(B)$ in B. Suppose B has one minimal delinquent chain initiated by c_0 and that the creation steps below $dels(B)$ in B all have leading ones in distinct places, all of which are strictly larger than the location of the leading one for c_0. Then S applies to B.*

PROOF. For each internal m, there must be a c below $dels(B)$ precluding the downward shifting of m, implying we do have at least one such c. In order for c_0 to be the first label of a delinquent chain, it must be reducible using creation steps from below it in the chain. If all those of B have larger leading one than c_0, then they cannot be used in the reduction of c_0 unless they are themselves first reduced. But if some such c is reducible not using c_0, then this enables the application of the splitting map to B.

Proposition 8.32. *Suppose some creation step c_1 below $dels(B)$ in B has leading one in the same place as c_0. Then B has no internal merge steps, except perhaps ones merging c_0 with vector spaces created prior to B.*

PROOF. Letting $c_1, c_2, \ldots c_r$ be the creation steps below $dels(B)$ in B with $c_1 < c_2 < \cdots < c_r$, then for the labels of $dels(B)$ to be ascending, we must merge c_1 with c_2 before merging in any other of these spaces.

In fact, $m(c_1, c_2)$ will be the earliest merge step in $dels(B)$ unless there are vector spaces created below B with even smaller vectors than c_1 that are also getting merged. But these steps cannot reduce c_0. Therefore, they are shiftable to below $dels(B)$, unless they merge c_0 with a vector space created in a lower λ-block.

Define $im(S^k)$ recursively by $(B_1, \dots, B_k) \in im(S^{k-1})$ for $(B_1, B_2) \in im(S)$ and $(B_1 \cup B_2, B_3, \dots, B_k) \in im(S^{k-2})$.

Lemma 8.33. *$(B, B_1, \dots, B_i) \in im(S^i)$ for $i \geq 2$ implies $(B_1, \dots, B_i) \in im(S^{i-1})$.*

PROOF. We begin with the case $i = 2$. Consider $(B, B_1, B_2) \in im(S^2)$. Let $B \cup B_1$ denote $S^{-1}(B, B_1)$, and let $B \cup B_1 \cup B_2$ denote $S^{-1}(B \cup B_1, B_2)$. Let c_0, c_1, c_2 initiate $dels(B), dels(B_1), dels(B_2)$, respectively. Either $dels(B \cup B_1 \cup B_2)$ has one delinquent chain initiated by c_2, or it has two delinquent chains, the higher of which is initiated by c_2. Either way, $B \cup B_1$ is the optimal block to split off from $B \cup B_1 \cup B_2$, i.e. is the lower block created by S. Therefore B_1 is likewise split off from $B_1 \cup B_2$ by S when $B_1 \cup B_2$ is preceded by B, since splitting B off first only impacts the optimality choice in that it ensures that all labels created in B have already been created prior to making this choice. Since S preserves the relative order of c_0, c_1, c_2,we may iterate this argument for $i > 2$. □

Lemma 8.34. *Consider maximal chains which coincide on all but an interval where they all have the label set $B_1 \cup B_2 \cup B_3$, and for shorthand denote these maximal chains by just their set of λ-blocks on this interval If (B_1, B_2, B_3) has $(B_1, B_2) \in im(S)$ and $(B_2, B_3) \in im(S)$, then $B_1 \cup B_2 \cup B_3$ is matched with either $(B_1 \cup B_2, B_3)$ or $(B_1, B_2 \cup B_3)$, except in the case where (B_1, B_2) and (B_2, B_3) are both in $im(S_2)$.*

PROOF. We break this into three cases, based on whether (B_1, B_2) and (B_2, B_3) are in the image of S_1 or S_2. If both are in the image of S_1, then $S_1(B_1 \cup B_2 \cup B_3) = (B_1 \cup B_2, B_3)$ follows from the fact that $S_1(B_2 \cup B_3) = (B_2, B_3)$ for $B_2 \cup B_3$ appearing immediately above λ-block B_1; the point is that the label c_i initiating $dels(B_2)$ should also initiate the lowest minimal delinquent chain in the lower λ-block obtained by applying S_1 to $B_1 \cup B_2 \cup B_3$, that the creation step c_j which is the largest pulled into the lower λ-block will be the same as in the lower λ-block of $S_1(B_2, B_3)$ and that those creation steps in $B_1 \cup B_2$ must be exactly the creation steps c having $m(c) \leq m(c_j)$, implying that exactly these will be pulled into the lower λ-block. From this the case follows. For $(B_1, B_2) \in im(S_2)$ and $(B_2, B_3) \in im(S_1)$

we get $S_1(B_1 \cup B_2 \cup B_3) = (B_1, B_2 \cup B_3)$, again yielding the desired result. Finally, for $(B_1, B_2) \in im(S_1)$ and $(B_2, B_3) \in im(S_2)$ we get $S_2(B_1 \cup B_2 \cup B_3) = (B_1 \cup B_2, B_3)$, still once again yielding the result in this case. $\square$

Corollary 8.35. *Any critical cell with a λ-block in the preimage of the splitting map will be matched in round two. Any pair $(B, B') \in im(S)$ will be matched in round two, unless it is of a type to be matched in round four.*

8.8. **Critical cell matching: round three.** This section deals with merge steps m in the interior of a delinquent chain in a maximal chain contributing a critical cell not matched in round one or two. Specifically, we must have $dels(B)$ initiated by c_0 with exactly one creation step c_1 below $dels(B)$ in B, and we must have that c_0, c_1 have leading ones in exactly the same place. The only possible internal merge steps are ones merging c_0 with spaces created prior to B.

Matching 8.36. *In this case, we match by shifting the largest such internal merge step m to immediately above $dels(B)$, i.e. into the descending portion of the λ-block immediately above B. It is possible this critical cell could have already been matched by shifting this m further upward into the delinquent chain of this λ-block, but if so we may propagate this process upward redistributing the matching.*

Proposition 8.37. *Every remaining critical cell σ as above. is matched bijectively with a critical cell τ as above.*

PROOF. The idea is to match σ by shifting m upward to just above the delinquent chain. If the resulting cell has already been matched by shifting m upward still farther into the next delinquent chain, then we may keep repeating this process of upward shifting of m. This process will eventually terminate, since we have the option of shifting m to above the highest delinquent chain overall.

In effect, this reverses a gradient path in P^M (not to be confused with a covering relation in P^M) in order to match two previously unmatched critical cells.

Proposition 8.38. *The third round of matching is acyclic.*

PROOF. Suppose there were a directed cycle C in the multi-graph face poset $P^{M''}$ of critical cells which survived the first two rounds of cancellation. Consider the lowest internal merge step m seen anywhere in this cycle. From this point, the cycle must have an upward edge shifting this m to a higher position; however, the cycle edge leading into this cycle position must be a downward edge moving m into its internal position, and the only possibility is that this downward edge is exactly the same covering relation as the upward edge, a contradiction. $\square$

8.9. **Critical cell matching: round four, the final round.** We have two remaining types of λ-blocks B still requiring cancellation:

(1) B with two overlapping delinquent chains that would have been matched in round one by shifting some $c \in B$ upward into $dels(B)$, except that this would create a third overlapping minimal delinquent chain initiated by c, causing the J-intervals not to cover all ranks, so that the maximal chain does not contribute a critical cell
(2) B with two overlapping delinquent chains to which S_2 applies yielding (B_1, B_2), with B immediately preceded by some B' with $(B', B_1) \in im(S_2)$.

The latter case was not covered by the parity argument in Section 8.6 because $(B', B_1, B_2) \notin im(S_2^2)$; this is because maximal chains with three or more overlapping minimal delinquent chains do not contribute critical cells.

We will match and cancel all of the critical cells of type one with all of the critical cells of type two. First we give an example that captures the idea: consider the label sequence $c_1, c_0, c_2, m(c_0), m(c_1), m(c_2)$ which has two overlapping delinquent chains, i.e. a λ-block of type (1). This is matched with the type (2) λ-block labelled $c_0, m(c_0), c_1, c_2, m(c_1), m(c_2)$. The gradient path from the latter to the former first swaps $m(c_0), c_1$, then has an upward step swapping c_1, c_0, then has downward step swapping $m(c_0), c_2$. The former cell is not yet matched because its unanchored creation step cannot shift to the interior of $dels(B)$ to get a critical cell, since then the J-intervals would not cover all ranks. The latter cell was not matched by applying the splitting map to $c_1, c_2, m(c_1), m(c_2)$ to obtain $c_1, m(c_1), c_2, m(c_2)$ because the latter is instead matched by unsplitting the pair of blocks $c_0, m(c_0), c_1, m(c_1)$.

Proposition 8.39. *All critical cells of types (1) and (2) may be matched and cancelled simultaneously by an acyclic matching.*

PROOF. First note that all unmatched critical cells with any λ-blocks of type (1) or (2) above may be matched with one of the other type by a gradient path just like the one described above. The necessary conditions for type (1) imply those for type (2) and conversely. Moreover there is always a gradient path similar to above, and it is always the unique gradient path between these two unmatched cells. This is clearly a bijection between types (1) and (2), when one classifies cells based on the lowest λ block of either type. There are no extraneous labels, since those would have caused the cell to have been matched earlier.

To verify acyclicity, notice that each matching step, reduces by one the number of creation steps just below λ-blocks with two overlapping delinquent chains. Thus, each downward step must increase by one the number of such creation steps. Also each upward step splits a λ-block, so each downward step must merge a pair of λ blocks. However, it is impossible for a downward step to do both these things unless we are traversing a matching edge between types (1) and (2). In this case it has the other orientation, unless it is at higher ranks than the upward edge just traversed, making it impossible to ever complete a cycle. □

Finally, we have matched enough cells to conclude:

Theorem 8.40. *The order complex of $PD_n(q)$ is homotopy equivalent to a wedge of spheres of top dimension.*

PROOF. We have constructed an acyclic matching on P^M which leaves only top-dimensional critical cells unmatched. By Theorem 6.7, the gradient paths given by matched critical cells in P^M may be reversed simultaneously to obtain a discrete Morse function which only has top-dimensional critical cells. This implies the result, as noted in Remark 6.1, it allows us to collapse $\Delta(PD_n(q) \setminus \{\hat{0}, \hat{1}\})$ onto a homotopy equivalent CW-complex with only top-dimensional cells and a single base point. □

9. The Cohen-Macaulay property for $PD_n(q)$

This section generalizes Theorem 8.40 to arbitrary poset intervals.

Proposition 9.1. *If each interval $(x, \hat{1})$ has the homotopy type of a wedge of spheres of top dimension, then $PD_n(q)$ is homotopically Cohen-Macaulay.*

PROOF. Notice that the interval $(\hat{0}, v)$ is a product of posets $PD_m(q)$ for various $m < n$, one for each component of v. More generally, (u, v) is a product of intervals of the form $(x_i, \hat{1})$ within these smaller posets $PD_m(q)$. Thus, we may show that each interval of the form $(x_i, \hat{1})$ is homotopy-equivalent to a wedge of spheres of top dimension to conclude that the product must also be, since the homotopy type of a product of posets is a join of the homotopy types of the individual posets. □

Given a matroid M with ground set E and an independent set I, recall that the *contraction* of M by I, denoted M/I, is a matroid with ground set $E \setminus I$ whose independent sets are sets $J \subseteq E \setminus I$ such that $J \cup I$ is independent in M.

Theorem 9.2. *For any $x \in PD_n(q)$, $\Delta(x, \hat{1})$ is homotopy equivalent to a wedge of spheres of top dimension.*

PROOF. Let B_x be a basis for the collective span of the components of x. We may split $(x, \hat{1})$ into cell-blocks by letting the frames be the set of bases for the matroid contraction $M \setminus B_x$. Our arguments from earlier sections again show that we have an EL-labelling for each cell-block, that the facet order based on ordering frames and using this EL-labelling within each cell-block gives a well-behaved Morse function, and that all non-top-dimensional critical cells may be cancelled simultaneously.

The point is that saturated chains from x to $\hat{1}$ will again have a λ-block structure whenever they contribute a critical cell, and that non-top-dimensional critical cells will imply delinquent chains with internal merge or creation steps; these are matched just as before, noting that all matching steps will allow us to preserve the chain element x, thereby matching with a partner cell also in the interval. $\square$

Combining Proposition 9.1 and Theorem 9.2 yields:

Corollary 9.3. *$PD_n(q)$ is homotopically Cohen-Macaulay. Hence, any rank-selected subposet is also Cohen-Macaulay.*

Munkres proved in [Mu, Cor. 6.6] that every rank-selection of a Cohen-Macaulay poset is Cohen-Macaulay. Thus, rank-selections of $PD_n(q)$ have homology concentrated in top degree, suggesting there could be interesting representations of $GL_n(q)$ on top homology of rank-selected subposets, much like the symmetric group representations for various rank-selections in the partition lattice.

Question 9.4. *The character formula provided in the next section will yield*

$$\mu_{PD_n(q)}(\hat{0}, \hat{1}) = \frac{-1}{n} q^{\binom{n}{2}} \prod_{j=1}^{n-1} (q^j - 1) = -\frac{|GL_n(q)|}{n \cdot (q^n - 1)}$$

as the dimension of the $GL_n(q)$-representation on top homology multiplied by a suitable sign. Is there a nice description of surviving critical cells that gives an alternate explanation for this formula? Is there a nice formula for $\mu_{PD_n(q)}(x, y)$ for arbitrary $x, y \in PD(n, q)$?

Question 9.5. *How is the Möbius function for $PD_n(q)$ related to the number of bases with internal activity zero in the matroid of lines in a finite vector space? This latter quantity was expressed as an alternating sum in [KRS].*

10. The proof of Theorem 1.2

For a positive integer n and a prime power q let $V = V(n, q)$ be an n-dimensional vector space over $\mathbb{F}_q$. The lattice of subspaces of V is denoted by $B_n(q)$. The action of $GL(V) = GL_n(q)$ on V determines an action of $GL_n(q)$ on $B_n(q)$. As noted in the introduction, R. Steinberg showed in [Ste] that $(-1)^n \Lambda_{B_n(q),GL_n(q)}$ is an irreducible character (which we call St) of $GL_n(q)$. This character, which is known as the *Steinberg character* of $GL_n(q)$, plays an important role in the representation theory of $GL_n(q)$ (see for example [Ca, Chapter 6]).

We prove Theorem 1.2 by induction on n, the case $n = 1$ being easy since $PD_1(q) = \{\hat{0}, \hat{1}\}$ and $GL_1(q)$ is a Coxeter torus in itself.

Considering the Whitney homology of the Cohen-Macaulay poset $PD_n(q)$ (see for example [Su1, Section 1]), we obtain the identity

$$\sum_{X \in PD_n(q)} (-1)^{rk(X)} \widetilde{H}_{rk(X)-2}(\Delta(\hat{0}, x)) = 0 \tag{1}$$

of virtual $GL_n(q)$-modules. For $X = \{X_1, \dots, X_k\} \in PD_n(q)$, let $\overline{X} = \bigoplus_{i=1}^k X_i$ and set

$$SD_n(q) := \{X \in PD_n(q) : \overline{X} = \mathbb{F}_q^n\}$$

and

$$RD_n(q) := PD_n(q) \setminus SD_n(q).$$

Using the Quillen Fiber Lemma [Qu], we see that the ($GL_n(q)$-equivariant) map from $RD_n(q) \setminus \{\hat{0}\}$ to the proper part of $B_n(q)$ given by $X \mapsto \overline{X}$ induces a homotopy equivalence of order complexes. Therefore, the homology of $\Delta(RD_n(q) \setminus \{\hat{0}\})$ is concentrated in dimension $n - 2$ and the unique nontrivial homology group provides the Steinberg representation of $GL_n(q)$. Since $RD_n(q)$ is an ideal in $PD_n(q)$, the homology of the order complex of any interval $(\hat{0}, X)$ in $RD_n(q)$ is concentrated in dimension $rk(X) - 2$. Considering the Whitney homology of $RD_n(q) \cup \{\hat{1}\}$, we obtain the identity

$$(-1)^n St + \sum_{X \in RD_n(q)} (-1)^{rk(X)} \widetilde{H}_{rk(X)-2}(\Delta(\hat{0}, x)) = 0 \tag{2}$$

of virtual $GL_n(q)$-modules. Combining identities (1) and (2), we obtain the identity

$$\sum_{X\in SD_n(q)} (-1)^{rk(X)} \widetilde{H}_{rk(X)-2}(\Delta(\hat{0},x)) = (-1)^n St \tag{3}$$

of virtual $GL_n(q)$-modules.

Let $X = \{X_1, \dots, X_k\} \in SD_n(q)$. Set $d_i = \dim X_i$. We may assume that $d_i \geq d_{i+1}$ for all $i \in [k-1]$. Then $\lambda(X) := (d_1, \dots, d_k)$ is a partition of n. For each partition $\lambda = (\lambda_1, \dots, \lambda_k)$ of n, set

$$SD_\lambda(q) := \{X \in SD_n(q) : \lambda(X) = \lambda\}.$$

For every $X \in SD_\lambda(q)$, we have

$$(\hat{0}, X) \cong \prod_{i=1}^{k} PD_{\lambda_i}(q),$$

so $rk(X) = 2n - k$. Also, $GL_n(q)$ acts transitively on $SD_\lambda(q)$, and it follows that if G_λ is the stabilizer of any $X_\lambda \in SD_\lambda(q)$ then we have the identity

$$\sum_{X\in SD_\lambda(q)} \widetilde{H}_{2n-k-2}(\Delta(\hat{0},X)) = \widetilde{H}_{2n-k-2}(\Delta(\hat{0},X_\lambda))_{G_\lambda}^{GL_n(q)} \tag{4}$$

of $GL_n(q)$-modules. (Here we are using the usual notation for induced modules.)

Fix a partition λ of n and for $j \in [n]$ let $m_j(\lambda)$ be the number of parts of λ equal to j. Then

$$G_\lambda \cong \prod_{m_j(\lambda)\neq 0} S_{m_j}[GL_j(q)].$$

Indeed, G_λ is the stabilizer in $GL_n(q)$ of some $X_\lambda \in SD_n(q)$ which contains $m_j := m_j(\leq)$ subspaces of dimension j for each j. The factor $G_\lambda^{(j)} := S_{m_j}[GL_j(q)]$ in the given direct product is a wreath product in which the kernel $GL_j(q)^{m_j}$ acts componentwise on the sum of these m_j subspaces and a complement S_{m_j} permutes them. This factor acts trivially on the span of the subspaces in X_λ that do not have dimension j.

For $j \in [n]$, set

$$X_\lambda^{(j)} := \{W \in X_\lambda : \dim(W) = j\}.$$

Then

$$(\hat{0}, X_\lambda) \cong \prod_{m_j\neq 0} (\hat{0}, X_\lambda^{(j)}). \tag{5}$$

The group G_λ acts on $\prod_{m_j\neq 0}(\hat{0}, X_\lambda^{(j)})$. Each factor $G_\lambda^{(j)}$ acts trivially on $(\hat{0}, X_\lambda^{(i)})$ for $i \neq j$ and acts on $(\hat{0}, X_\lambda^{(j)})$ according its action on $X^{(j)}$. Some isomorphism giving (5) is G_λ-equivariant. It follows (see [Su, Proposition 2.1]) that we have an isomorphism of G_λ-modules

$$\widetilde{H}_{rk(X_\lambda)-2}(\Delta(\hat{0}, X_\lambda)) \cong \bigotimes_{m_j\neq 0} \widetilde{H}_{rk(X_\lambda^{(j)})-2}(\Delta(\hat{0}, X_\lambda^{(j)})). \tag{6}$$

If $m_j > 0$, there is a $G_\lambda^{(j)}$-equivariant isomorphism between the interval $(\hat{0}, X_\lambda^{(j)})$ and the direct product $PD_j(q)^{m_j}$. (The group $G_\lambda^{(j)}$ acts on the direct product with the kernel $GL_j(q)^{m_j}$ acting componentwise and a complement S_{m_j} permuting the components.) If V_j is a 1-dimensional space on which the complement S_{m_j} acts by the sign representation and the kernel acts trivially, then (see [Su, Proposition 2.3]) we have an isomorphism of $G_\lambda^{(j)}$-modules

$$\widetilde{H}_{rk(X_\lambda^{(j)})-2}(\Delta(\hat{0}, X_\lambda^{(j)})) \cong V_j \otimes \bigotimes^{m_j} \widetilde{H}_{2j-3}(\Delta(PD_j(q))), \tag{7}$$

where the complement permutes the m_j tensor factors $\widetilde{H}_{2j-3}(\Delta(PD_j(q)))$ and the kernel acts componentwise.

Assume now that $\lambda \neq (n)$. For $j \in \mathbb{N}$, let N_j be the normalizer of a Coxeter torus T_j in $GL_j(q)$. For $j < n$, our inductive hypothesis says that there is an isomorphism

$$\widetilde{H}_{2j-3}(\Delta(PD_j(q))) \cong \Theta_j = (\theta_j)_{N_j}^{GL_j(q)},$$

where θ_j is a linear character of N_j that is trivial on T_j and faithful on a complement to T_j in N_j. It now follows from (7) that

$$\widetilde{H}_{rk(X_\lambda^{(j)})-2}(\Delta(\hat{0}, X_\lambda^{(j)})) \cong \left(V_j \otimes \bigotimes^{m_j} \mathbb{C}\right)_{S_{m_j}[N_j]}^{G_\lambda^{(j)}}, \tag{8}$$

where the kernel $N_j^{m_j}$ of $S_{m_j}[N_j]$ acts componentwise according to θ_j on $\bigotimes^{m_j}\mathbb{C}$ and trivially on V_j, while a complement S_{m_j} acts by the sign representation on V_j and permutes the components of $\bigotimes^{m_j}\mathbb{C}$. (Indeed, let V be any S_m-module, Let H be a finite group, let $K \leq H$ and let W be a K-module. Then the $S_m[H]$-modules $V \otimes \bigotimes^m W_K^H$ and $(V \otimes \bigotimes^m W)_{S_m[K]}^{S_m[H]}$ are isomorphic - we postpone our elementary proof of this fact.)

Combining (6) and (8), we get

$$\widetilde{H}_{rk(X_\lambda)-2}(\Delta(\hat{0}, X_\lambda)) \cong \left(\bigotimes_{m_j \neq 0} \left(V_j \otimes \bigotimes^{m_j} \mathbb{C} \right) \right)^{G_\lambda}_{\prod_{m_j \neq 0} S_{m_j}[N_j]}. \tag{9}$$

(It is well known and not hard to prove that if $K_i \leq H_i$ for $i \in [k]$ and W_i is a K_i-module then the $\prod_i H_i$-modules $(\bigotimes_i W_i)^{\prod_i H_i}_{\prod_i K_i}$ and $\bigotimes_i (W_i)^{H_i}_{K_i}$ are isomorphic.)

For each partition λ of n, there is a maximal torus $T_\lambda \leq GL_n(q)$ such that $T_\lambda \cong \prod_{i=1}^{l(\lambda)} \mathbb{Z}_{q^{\lambda_i}-1}$ and the group $\prod_{m_j \neq 0} S_{m_j}[N_j]$ appearing in (9) is the normalizer N_λ of T_λ in $GL_n(q)$. (See for example [Ca, Chapter 3].) Following [Su], we will from now on write $A_{m_j}[\theta_j]$ for the character $(V_j \otimes \bigotimes^{m_j} \mathbb{C})$ of $S_{m_j}[N_j]$ appearing in (8).

From (3) and (9), along with transitivity of induction, we get an isomorphism

(10)

$$-\widetilde{H}_{2n-3}(PD_n(q)) + \sum_{\lambda \neq (n)} (-1)^{l(\lambda)} \left(\bigotimes_{m_j(\lambda) \neq 0} A_{m_j(\lambda)}[\theta_j] \right)^{GL_n(q)}_{N_\lambda} = (-1)^n St$$

of virtual $GL_n(q)$-modules, where $l(\lambda)$ is the number of parts of a partition λ. (Note that $rk(X_\lambda) \equiv l(\lambda) \bmod 2$.)

In the notation from [Su] introduced above, we have $A_1[\theta_n] = \theta_n$. For each partition λ of n, we fix $g(\lambda) \in S_n$ of cycle shape λ. Let sgn be the sign character of S_n. Then

$$\mathrm{sgn}(g(\lambda)) = (-1)^{n-l(\lambda)}. \tag{11}$$

Combining (10) and (11), we get the following key result.

Proposition 10.1. *Theorem 1.2 holds if and only if there is an isomorphism*

$$St \cong \sum_{\lambda \vdash n} sgn(g(\lambda)) (\otimes_{j=1}^n A_{m_j(\lambda)}[\theta_j])^{GL_n(q)}_{N_\lambda} \tag{12}$$

of virtual $GL_n(q)$-modules.

Our goal is now to prove that identity (12) does indeed hold by applying the theorem of Srinivasan mentioned in the introduction. As noted in the introduction, there exists an isomorphism

$$f_\lambda : N_\lambda / T_\lambda \to C_{S_n}(g(\lambda)).$$

Since T_λ lies in the kernel of $(\otimes_{j=1}^{n} A_{m_j(\lambda)}[\theta_j])$, that character determines a character ρ_λ of N_λ/T_λ. We can now define a character ψ_λ of $C_{S_n}(g(\lambda))$ by

$$\psi_\lambda(k) := \rho_\lambda(f_\lambda^{-1}(k)).$$

Further, if $x \in S_n$ and $g = g(\lambda)^x$, we can define a character ψ_g of $C_{S_n}(g) = C_{S_n}(g(\lambda))^x$ by

$$\psi_g(k^x) = \psi_\lambda(k).$$

It is straightforward to show (using the fact that ψ_λ is a class function) that ψ_g depends only on g (that is, not on the choice of x), so we now have a well-defined character ψ_g of $C_{S_n}(g)$ for each $g \in S_n$ of cycle shape λ, and letting λ run through the set of all partitions of n we obtain a character ψ_g for each $g \in S_n$.

Now $C_{S_n}(g)$ acts on the set of cycles of g. Let $K(g)$ be the subgroup of $C_{S_n}(g)$ consisting of those elements which act trivially on this set. In our setting, Srinivasan's theorem says the following.

Theorem 10.2 (see [Sr]). *Identity (12) holds in the character ring of $GL_n(q)$ if and only if for all nonidentity $x \in S_n$ and all $c \in C_{S_n}(x)$ we have*

$$\sum_{g \in K(x)c} sgn(g)\psi_g(x) = 0. \tag{13}$$

It should be noted that as stated in [Sr], Srinivasan's theorem applies to a finite group which is the set of rational points of an endomorphism of a simply connected, semisimple algebraic group over an algebraically closed field of characteristic $p > 0$ such that p does not divide the order of the associated Weyl group. Now $GL_n(q)$ is the set of rational points of a Frobenius automorphism of a connected (but not simply connected), reductive (but not semisimple) algebraic group, and no mention has yet been made of the relation of the characteristic p to the order of the Weyl group S_n, so some further explanation is necessary. The theorem depends on nine lemmas stated in [Sr], the first seven of which apply to connected reductive groups independent of characteristic. The stronger assumptions are used in Lemma 8 to show that certain pairs of elements of the finite group which are conjugate in the algebraic group are already conjugate in the finite group. This property holds for all pairs of elements of $GL_n(q)$. In Lemma 9, simple connectivity is used to assert that the centralizer of a semisimple element of the algebraic group is a connected reductive group. This property holds for any semisimple element of a general linear group over an algebraically closed field of positive characteristic.

Let us now prove that the identity (13) holds. First, we must take a closer look at the groups $C_{S_n}(g)$ and the characters ψ_g. We consider first the case where $n = lm$, $1 \leq l \leq n$ and g has cycle shape $\mu = (l, l, ..., l)$. For the rest of the paper, we shall assume that when this case holds, we have identified $[n]$ with $\mathbb{Z}_l \times \mathbb{Z}_m$ in such a way that

$$(i, j)g = (i + 1, j)$$

for all $(i, j) \in \mathbb{Z}_l \times \mathbb{Z}_m$.

Then $K(g)$ consists of all $k \in S_n$ such that there exist $r_{1,k}, \ldots, r_{m,k} \in \mathbb{Z}_l$ such that

$$(i, j)k = (i + r_{j,k}, j)$$

for all $(i, j) \in \mathbb{Z}_l \times \mathbb{Z}_m$. Define $H(g)$ to be the set of all $h \in S_n$ for which there exists some $\tau_h \in S_m$ such that

$$(i, j)h = (i, j\tau_h)$$

for all $(i, j) \in \mathbb{Z}_l \times \mathbb{Z}_m$. Then $H(g)$ is a complement in $C_{S_n}(g)$ to the normal subgroup $K(g)$. Moreover, we may (and do from now on) assume that f_μ has been chosen so that for $h \in H(g)$ and $k \in K(g)$ we have

$$\psi_g(hk) = sgn(\tau_h)e^{2\pi i \sum_{j=1}^{m} r_{j,k}/l}. \tag{14}$$

Now for arbitrary $g \in S_n$ having cycle shape λ, write $g = g_1 \ldots g_n$, where g_j is a product of $m_j(\lambda)$ j-cycles (and the g_j have disjoint supports). Let X_1 be the set of fixed points of g and, for $j > 1$, let X_j be the support of g_j. We have

$$C_{S_n}(g) = \prod_{j=1}^{n} C_{S_{X_j}}(g_j),$$

and we may (and do from now on) assume that f_λ has been chosen so that for $c = (c_1, \ldots, c_n) \in C_{S_n}(g)$ (with $c_j \in C_{S_{X_j}}(g_j)$), we have

$$\psi_g(c) = \prod_{j=1}^{n} \psi_{g_j}(c_j). \tag{15}$$

From (14) and (15), we get the following result.

Lemma 10.3. *Let $g_1, g_2 \in S_n$ have disjoint supports. For $i = 1, 2$ let $x_i \in C_{S_n}(g_i)$ with the support of x_i contained in that of g_i. Then*

$$\psi_{g_1 g_2}(x_1 x_2) = \psi_{g_1}(x_1)\psi_{g_2}(x_2).$$

We will use Lemma 10.3 to reduce the claim that identity (13) holds for all $g \in S_n$ and all $c \in C_{S_n}(g)$ to the claim that (13) holds in certain special cases.

We return now to the case where g has shape $\mu = (l, \ldots, l)$. Let $c = hk \in C_{S_n}(g)$ with $h \in H(g)$ and $k \in K(g)$. We call c a *standard m-cycle on g* if τ_h is an m-cycle in S_m and $\sum_{j=1}^{m} r_{j,k} \equiv 0$. (Note that in this case c has cycle shape $(m, \ldots, m)$.

Lemma 10.4. *To prove that identity (13) holds for all $1 \neq g \in S_n$ and all $c \in C_{S_n}(g)$, it suffice to prove that (13) holds when $n = lm$ ($l > 1$), g has shape $(l, \ldots, l)$ and c is a standard m-cycle on g.*

PROOF. Fix $g \in S_n$ and $c \in C_{S_n}(g)$. Write $g = g_1 \ldots g_r$ and $c = c_1 \ldots c_r$ so that the orbits of $\langle g, c\rangle$ are the supports of the groups $\langle g_i, c_i\rangle$, $i \in [r]$. (If $|X| = 1$ then the support of S_X is defined to be X.) For each $i \in [r]$, g_i has shape $l_i^{m_i}$ and c_i acts as an m_i-cycle on the set of cycles of g_i. Let X_i be the support of $\langle g_i, c_i\rangle$ and set $C_i = C_{S_{X_i}}(g_i)$. Let K_i be the subgroup of C_i consisting of those elements which act trivially on the set of cycles of g_i. Each $k \in K(g)$ can be written (uniquely) as $k = k_1 \ldots k_r$ with $k_i \in K_i$. By Lemma 10.3 (and the distributive law), we have

$$\sum_{x \in K(g)c} sgn(x)\psi_x(g) = \prod_{i=1}^{r} \sum_{y \in K_i c_i} sgn(y)\psi_y(g_i).$$

Since c_i acts as an m_i-cycle on the set of cycles of g_i, we see that there is some $k_i \in K_i$ such that $k_i c_i$ is a standard m_i-cycle on g_i (in S_{X_i}). Moreover, $K_i k_i c_i = K_i c_i$ and the lemma follows. □

Lemma 10.5. *Let $n = lm$, and let $g \in S_n$ have cycle shape $(l, \ldots, l)$. Then all standard m-cycles on g are conjugate in $C_{S_n}(g)$.*

PROOF. Let $c = hk$ be a standard m-cycle on g with $h \in H(g)$, $k \in K(g)$, so

$$(i, j)c = (i + r_{j\tau_h,k}, j\tau_h)$$

for all $(i, j) \in \mathbb{Z}_l \times \mathbb{Z}_m$. The projection of $C_{S_n}(g)$ on $H(g)$ followed by the isomorphism $x \mapsto \tau_x$ is surjective onto S_m. Since all m-cycles in S_m are conjugate, we may assume that $j\tau_h = j + 1$ for all $j \in \mathbb{Z}_m$. Thus

$$(i, j)c = (i + r_{j+1,k}, j + 1).$$

Let $x \in K(c)$ satisfy

$$r_{j,x} - r_{j+1,x} \equiv r_{j+1,k}$$

for all $j \in \mathbb{Z}_m$. (These equations can be solved simultaneously, since $\sum_{j=1}^{m} r_{j,k} \equiv 0$.) Then

$$(i,j)x^{-1}cx = (i, j+1)$$

for all $(i,j) \in \mathbb{Z}_l \times \mathbb{Z}_m$. Thus every standard m-cycle on g is conjugate to the unique $h \in H(g)$ such that $\tau_h = (1 \dots m)$. $\square$

We are now ready to complete our proof that identity (13) holds.

Proposition 10.6. *Let $n = lm$ $(l > 1)$ and let $g \in S_n$ have cycle shape $(l, \dots, l)$. Let c be a standard m-cycle on g. Then*

$$\sum_{x \in K(g)c} sgn(x)\psi_x(g) = 0.$$

PROOF. Note that if x, y are conjugate in $C_{S_n}(g)$, then $\mathrm{sgn}(x) = \mathrm{sgn}(y)$ and $\psi_x(g) = \psi_y(g)$. Since $y^{-1}K(g)cy = K(g)y^{-1}cy$ for all $y \in C_{S_n}(g)$, we may (and do) assume by Lemma 10.5 that

$$(i,j)g = (i+1, j)$$

and

$$(i,j)c = (i, j+1)$$

for all $(i,j) \in \mathbb{Z}_l \times \mathbb{Z}_m$. Now let $k \in K(g)$. Continuing with the notation fixed above, we have

$$(i,j)kc = (i + r_{j,k}, j+1)$$

for all $(i,j) \in \mathbb{Z}_l \times \mathbb{Z}_n$. Set

$$r = r(k) := \sum_{j=1}^{m} r_{j,k}.$$

Let $t = t(r) = \gcd(l, r)$ and set $d = l/t$. Then $\langle kc \rangle$ has exactly t orbits $O_1, \dots, O_t$ on $\mathbb{Z}_l$, with $(i,1) \in O_i$ for each $i \in [t]$. Moreover, $|O_i| = md$ for each i. Define $h, b \in C_{S_n}(kc)$ by

$$(i,j)h = \begin{cases} (i+1, j) & i \notin O_t, \\ (i+1-t, j) & i \in O_t, \end{cases}$$

and

$$(i,j)b = \begin{cases} (i,j) & i \notin O_1, \\ (i+t, j) & i \in O_1. \end{cases}$$

Then $g = hb$, h is a standard t-cycle on kc and $b \in K(kc)$. In fact, if we let c_1 be the cycle of kc whose support is O_1, then b is a power of c_1. Moreover, if $v = v(r) \in \mathbb{Z}_l$ satisfies $vr \equiv t$ in $\mathbb{Z}_l$ (and we abuse notation by letting v also stand for the smallest positive representative of $v \in \mathbb{Z}_l$) then we have $b = c_1^{mv}$. Collecting what we have learned so far, we have

$$\begin{aligned} sgn(kc)\psi_{kc}(g) &= (-1)^{t(md-1)}(-1)^{t-1}\psi_{c_1}(b) \\ &= (-1)^{n-1}e^{2\pi imv/md} \\ &= (-1)^{n-1}e^{2\pi ivt/l}. \end{aligned}$$

Now as k runs through K, $r(k)$ runs through $[l]$ l^{m-1} times. Since v, t depend only on r, it now suffices to show that

$$\sum_{r=1}^{l} e^{2\pi iv(r)t(r)/l} = 0 \tag{16}$$

We will show that (16) holds by showing that for each $j \in \mathbb{Z}_l$ there exists some (and therefore exactly one) $r \in \mathbb{Z}_l$ such that $v(r)t(r) \equiv j$ in $\mathbb{Z}_l$. Indeed, given j, let $x = \gcd(j, l)$. In $\mathbb{Z}_{l/x}$, j/x is a unit. Let $w \in \mathbb{Z}_l$ represent $(j/x)^{-1}$ in $\mathbb{Z}_{l/x}$ and let $r = wx \in \mathbb{Z}_l$. Then $t(r) = \gcd(r, l) = x$ (since $\gcd(w, l/x) = 1$). Moreover, $jw \equiv x \bmod l$ implies that $v(r) = j/x$. Thus $v(r)t(r) \equiv j$ as desired. □

11. $PD_n(q)$ as q-analogue of $\Pi_{\leq n}$

In this section we examine relations between the zeta polynomial, rank-generating function and characteristic polynomial for $PD_n(q)$ and those for $\Pi_{\leq n}$, showing that in the appropriate sense, $PD_n(q)$ is a q-analogue of the lattice of partial partitions. Let

$$\chi_P(t) = \sum_{u \in P} \mu_P(\hat{0}, u) t^{\mathrm{rk(P)-rk(u)}}$$

denote the characteristic polynomial of a ranked poset P. The zeta polynomial, denoted $Z_P(s)$, counts multichains of length s in P (and is shown in [St2] to extend to a polynomial in s). The rank generating function of P is the polynomial in variable x such that the coefficient of x^r is the number of elements of rank r in P.

For a real number q, let $[n]_q$ denote the q-analogue of n. That is, $[n]_q = 1 + q + \cdots + q^{n-1}$. The q-analogue of $\binom{n}{k}$, denoted $\left[{n \atop k}\right]_q$, is well-known to count k-dimensional subspaces of an n-dimensional vector space over a field of order q (e.g. see [St3]). It is a polynomial in q, in which the coefficient of q^r counts lattice paths from $(0, 0)$ to $(k, n-k)$ which consist of steps $(1, 0)$ and $(0, 1)$ and have area r below the path, within the box bounded by $(0, 0), (k, 0), (0, n-k)$, and $(k, n-k)$. For

$q \neq -1$, $\left[{n \atop k}\right]_q$ coincides with the rational function

$$\frac{[n]_q \cdots [1]_q}{([k]_q \cdots [1]_q)([n-k]_q \cdots [1]_q)}.$$

Lemma 11.1. *Fix $n \in \mathbb{N}$. Let $m \leq n$ and $\lambda = (\lambda_1, \ldots, \lambda_k)$ be a partition of n. There exists a polynomial $M_\lambda(t) \in \mathbb{Q}[t]$ (which does not depend on n) such that for each prime power q and each $w = \{W_1, \ldots, W_k\} \in PD_n(q)$ with $\dim(W_i) = \lambda_i$ for all i, we have $\mu_{PD_n(q)}(\hat{0}, w) = M_\lambda(q)$. Moreover, if $\pi = P_1| \ldots |P_k \in \Pi_{\leq n}$ with $|P_i| = \lambda_i$ for all i then $M_\lambda(1) = \mu_{\Pi_{\leq n}}(\hat{0}, \pi)$.*

PROOF. By the multiplicativity of the Möbius function on products of posets, it suffices to prove the claim when $\lambda = (n)$. By Corollary 1.3 we have

$$\mu_{PD_n(q)}(\hat{0}, \hat{1}) = -\frac{1}{n} \cdot q^{\binom{n}{2}} \prod_{i=1}^{n-1} (q^i - 1).$$

If we substitute $q = 1$ into the right hand side of the equation above, we get 0 unless $n = 1$, in which case we get -1. □

Lemma 11.2. *Fix n, m, λ as in Lemma 11.1. For a prime power q, let $f_{\lambda,n}(q)$ be the number of elements $\{W_1, \ldots, W_k\}$ in $PD_n(q)$ such that $\dim W_i = \lambda_i$ for all i and let $g(\lambda, n)$ be the number of elements $B_1| \cdots |B_k$ of $\Pi_{\leq n}$ in which $|B_i| = \lambda_i$ for all i. Then $f_{\lambda,n}$ is a polynomial in $\mathbb{Q}[q]$, and*

$$\lim_{x \to 1} f_{\lambda,n}(x) = g(\lambda, n).$$

PROOF. Set

$$h(\lambda, n, q) = \begin{bmatrix} n \\ \lambda_1 \end{bmatrix}_q \cdot \begin{bmatrix} n - \lambda_1 \\ \lambda_2 \end{bmatrix}_q \cdots \begin{bmatrix} n - (\lambda_1 + \lambda_2 + \cdots + \lambda_{k-1}) \\ \lambda_k \end{bmatrix}_q.$$

Let $m_i(\lambda)$ be the number of parts of λ of size i, and set $d_\lambda = \prod_{i \geq 1}(m_i(\lambda)!)$. Then $f_{\lambda,n}(q) = h(\lambda, n, q)/d_\lambda$ for all prime powers q, and

$$\lim_{x \to 1} \frac{h(\lambda, n, x)}{d_\lambda} = \frac{1}{d_\lambda} \cdot \binom{n}{\lambda_1, \ldots, \lambda_k} = g(\lambda, n).$$

□

Corollary 11.3. *Fix $n \in \mathbb{N}$. There is a function $R_n(t, x)$, polynomial in x and t, such that the rank generating function (with variable x) of $PD_n(q)$ is $R_n(q, x)$ for all prime powers q. Moreover, the rank generating function for $\Pi_{\leq n}$ is $\lim_{t \to 1} R_n(t, x)$.*

Proposition 11.4. *Fix $n \in \mathbb{N}$. There is a function $X_n(q,t)$, polynomial in q and t, such that $\chi_{PD_n(q)}(t) = X_n(q,t)$ for all prime powers q. Moreover,*

$$\lim_{q\to 1} X_n(q,t) = \chi_{\Pi_{\leq n}}(t).$$

PROOF. By Lemma 11.1, when calculating $\chi_{PD_n(q)}(t)$ we need only consider those elements of $PD_n(q)$ of the form $\{W_1,\dots,W_r\}$ where each W_i is 1-dimensional. On the other hand, we showed in Corollary 3.6 that the characteristic polynomial for $\Pi_{\leq n}$ may be computed by considering only those partial partitions in which each block has size one. Now the result follows easily from Lemma 11.2, since $\mu_P(\hat{0},u) = (-1)^{\mathrm{rk(u)}}$ for $P = PD_n(q)$ or $P = \Pi_{\leq n}$ and u any of the poset elements we are now summing over. □

Proposition 11.5. *Fix $n \in \mathbb{N}$. For each $s \in \mathbb{N}$, there is a polynomial $F_{n,s}(q)$ such that for all prime powers q, we have $Z_{PD_n(q)}(s) = F_{n,s}(q)$. Moreover,*

$$\lim_{q\to 1} F_{n,s}(q) = Z_{\Pi_{\leq n}}(s).$$

PROOF. The claim in the case $s = 1$ is an immediate consequence of Lemma 11.2, so we will prove the proposition by induction on s. Summing over all possible choices of the last element in a multichain of length s immediately yields

$$Z_{PD_n(q)}(s) = \sum_{u \in PD_n(q)} Z_{[\hat{0},u]}(s-1).$$

But $Z_{P\times Q}(s) = Z_P(s)\cdot Z_Q(s)$, so

$$\sum_{u\in PD_n(q)} Z_{[\hat{0},u]}(s-1) = \sum_{m\leq n}\sum_{\lambda\vdash m} f_{\lambda,n}(q) \prod_{i=1}^{l(\lambda)} Z_{PD_{\lambda_i}(q)}(s-1).$$

The result now follows from the inductive hypothesis. □

12. Potential generalizations to other matroids

In this section, we abstract some of the key ingredients from the homotopy type computation of earlier sections, in an effort to shed some light on what level of generality our homotopy type result might hold in (i.e. for which matroids). First is a property of circuits that is essentially Lemma R abstracted. To see this, notice that two lines with leading ones in the same place comprise a broken circuit.

Definition 12.1. *A matroid M has Property R if there exists an ordering on the ground set of M such that for any basis B and any pair v_i, v_j of internally passive elements of B, at least one of the following holds:*

(1) *v_i, v_j comprise a broken circuit*
(2) *$B \setminus \{v_j\}$ contains a v_i-broken circuit of M (as defined in Section 5)*
(3) *$B \setminus \{v_i\}$ contains a v_j-broken circuit of M*

The following may be useful for dealing with arbitrary poset intervals.

Proposition 12.2. *If a matroid M has Property R, then any contraction M/I also has Property R.*

PROOF. Let G be the ground set for M let $B = \{v_1, \ldots, v_r\}$ be a basis for M/I, and let v_i, v_j be internally passive in B. By definition of contraction, M has a basis $B' = \{w_1, \ldots, w_s, v_1, \ldots, v_r\}$ with $I = \{w_1, \ldots, w_s\}$. Clearly, v_i, v_j are internally passive in B' because the circuits in M/I which contain v_i, v_j as well as earlier elements of $G \setminus I$ result from circuits in M that also contain the same earlier elements of $G \setminus I$ to exchange for v_i, v_j, respectively. By Property R for M, either (1) v_i, v_j comprise a broken circuit in M, (2) v_i is internally passive in $B' \setminus \{v_j\}$ or (3) v_j is internally passive in $B' \setminus \{v_i\}$. (1) implies v_i, v_j also comprise a broken circuit in M/I. (2) implies v_i belongs to a circuit in $M \setminus \{v_j\}$ which includes exactly one element $w \in G \setminus B'$, and that $w < v_i$, so this contracts to a circuit in M/I which again allows v_i to be exchanged for w, since w is not in the flat spanned by I. Thus, v_i is internally passive in $B \setminus \{v_i\}$. Case (3) is similar to case (2). □

Remark 12.3. *Property R is not closed under deletion. For example, consider the graphic matroid with ground set $e_{1,3} < e_{1,2} < e_{2,3} < e_{3,4} < e_{1,4}$. This has Property R, but the matroid in which $e_{1,3}$ is deleted no longer has Property R.*

It is not hard to show that all graphic matroids resulting from threshold graphs have Property R. It seems likely that all graphic matroids resulting from chordal graphs have Property R.

Question 12.4. *Do all supersolvable matroids have Property R? Recall that the supersolvable graphic matroids are exactly those given by chordal graphs.*

Define the relation $\sim$ on the ground set for $PD_n(q)$ by $u \sim v$ whenever u, v have their leading ones in the same place.

Remark 12.5. $\sim$ *is an equivalence relation on the ground set for the matroid underlying* $PD_n(q)$*, and for its contractions.*

For $PD_n(q)$ as a whole, we have $u \sim v$ if and only if u, v comprise a broken circuit. For arbitrary intervals, $u \sim v$ implies u, v comprise a broken circuit, but the converse is not quite as clear.

One other key property (used in Lemma B) is as follows:

Definition 12.6. *A matroid has Property L if for any basis B and any $v_i, v_j, v_k \in B$ such that v_i, v_j are internally passive in B and $v_i < v_j < v_k$, then we have the following implication. If v_i is internally active in $B \setminus \{v_j\}$ and v_j is internally active in $B \setminus \{v_k\}$ then v_i is internally active in $B \setminus \{v_k\}$.*

Proposition 12.7. *All matroids satisfy Property L.*

PROOF. Any v_i-broken circuits C must also include v_j, which means C is also a v_j-broken circuit, since $v_j > v_i$. However, this means C also must include v_k. $\square$

Question 12.8. *Suppose a matroid M gives rise to a poset $PD(M)$ of partial decompositions into flats. If M has Property R, and M as well as all its contractions has an equivalence relation on its ground set based on broken circuits of size 2, then is $PD(M)$ Cohen-Macaulay?*

Notice that $PD(M)$ still has cell-blocks $\Pi_{\leq n}$, one for each basis of M, so perhaps our method of proof will generalize to this setting.

ACKNOWLEDGMENTS

The authors thank Richard Lyons, Vic Reiner, Bruce Sagan and Richard Stanley for helpful comments.

References

[BH] E. Babson and P. Hersh, *Discrete Morse functions from lexicographic orders*, Trans. Amer. Math. Soc. **357** (2005), no. 2, 509–534.

[Ba] H. Barcelo, *On the action of the symmetric group on the free Lie algebra and the partition lattice*, J. Combin. Theory Ser. A **55** (1990), 93–129.

[Be] A. Berglund, *Poincare' series of monomial rings*, preprint 2004.

[BBH] A. Berglund, J. Blasiak and P. Hersh, *Combinatorics of multigraded Poincare' series for monomial rings*, Preprint 2005.

[Bj1] A. Björner, *Shellable and Cohen-Macaulay partially ordered sets*, Trans. Amer. Math. Soc. **260** (1980), 159–183.

[Bj2] A. Björner, *Homology and shellability of matroids and geometric lattices*, in Matroid Applications (ed. N. White), Encyclopedia of Mathematics and Its Applications, **40** Cambridge Univ. Press, 1992.

[BW] A. Björner and M. Wachs, *On lexicographically shellable posets*, Trans. Amer. Math. Soc. **277, No. 1** (1983), 323–341.

[BWa] A. Björner and J. Walker, *A homotopy complementation formula for partially ordered sets*, Europ. J. Combin, **4** (1983), no. 1, 11–19.

[CHR] A.R. Calderbank, P. Hanlon and R.W. Robinson, *Partitions into even and odd block size*, Proc. London Math. Soc. (3), **53, No. 2** (1986), 288–320.

[Ca] R. Carter, *Finite groups of Lie type. Conjugacy classes and complex characters*, John Wiley and Sons, Inc., New York, 1985.

[Cr] H. Crapo, Möbius Inversion in Lattices, Arch. Math. (Basel) **19** (1968), 595-607.

[Ch] M. Chari, *On discrete Morse functions and combinatorial decompositions*, Discrete Math. **217** (2000), no. 1-3, 101–113.

[Do] Dowling, T.A., *A class of geometric lattices based on finite groups*, J. Combinatorial Theory Ser. B **14** (1973), 61–86.

[Fo] R. Forman, *Morse theory for cell complexes*, Adv. Math. **134** (1998), 90–145.

[Hal] P. Hall, *The Eulerian functions of a finite group*, Quart. J. Math. **7** (1936), 134–151.

[Ha1] P. Hanlon, *The fixed-point partition lattices*, Pacific J. Math. **96, No. 2** (1981), 319–341.

[Ha2] P. Hanlon, *A proof of a conjecture of Stanley concerning partitions of a set*, European J. Combin. **4** (1983), no. 2, 137–141.

[HH] P. Hanlon and P. Hersh, *Multiplicity of the trivial representation in rank-selected homology of the partition lattice*, J. Algebra **266** (2003), no. 2, 521–538.

[HaW] P. Hanlon and M. Wachs, *On Lie k-algebras*, Adv. Math. **113, No. 2** (1995), 206–236.

[He] P. Hersh, *On optimizing discrete Morse functions*, Adv. in Applied Math. **35** (2005), 294–322.

[HeW] P. Hersh and V. Welker, *Gröbner basis degree bounds on $Tor^{k[\Lambda]}(k,k)$ and discrete Morse theory for posets*, *Contemp. Math.*, **374**, American Mathematical Society, Providence, RI, 2005.

[HRW] P. Herzog, V. Reiner and V. Welker, *The Koszul property in affine semigroup rings*, Pacific J. Math, **186** (1998), 39–65.

[Is] I. M. Isaacs, *Character Theory of Finite Groups*, Pure and Applied Mathematics **69**, Academic Press, New York, 1976.

[Jo] J. Jonsson, *On the topology of simplicial complexes related to 3-connected and Hamiltonian graphs*, J. Combin. Theory, Ser. A. **104** (2003), no. 1, 169–199.

[Ke] A. Kerber, *Representations of permutation groups, II*, Lecture Notes in Mathematics **495**, Springer-Verlag, Berlin-New York, 1975.

[KRS] W. Kook, V. Reiner and D. Stanton, *Combinatorial Laplacians of matroid complexes*, J. Amer. Math. Soc., **13** (1999), no. 1, 129–148.

[Mc] P. McNamara, *EL-labelings, supersolvability and 0-Hecke algebra actions on posets*, J. Combin. Theory, Ser. A, **101** (2003), 69–89.

[Mu] J. Munkres, *Topological results in combinatorics*, Michigan Math. J., **31** (1984), 113–128.

[Ox] J. Oxley, *Matroid theory*, Oxford University Press, Oxford New York Tokyo, 1992.

[Qu] *Homotopy properties of the poset of nontrivial p-subgroups of a group*, Adv. in Math. **28** (1978), no. 2, 101–128.

[Sr] B. Srinivasan, On the Steinberg character of a finite simple group of Lie type, J. Austral. Math. Soc. **12** (1971), 1-14.

[St] R. Stanley, *Supersolvable lattices*, Algebra Universalis **2** (1972), 197–217.

[St1] R. Stanley, *Balanced Cohen-Macaulay complexes*, Trans. Amer. Math. Soc., **249** (1979), no. 1, 139–157.

[St2] R. Stanley, *Enumerative Combinatorics*, vol. 1, Wadsworth and Brooks/Cole, 1986.

[St3] R. Stanley, *Some aspects of groups acting on finite posets*, J. Combin. Theory Ser. A **32** (1982), no. 2, 132–161.

[Ste] R. Steinberg, *A geometric approach to the representations of the full linear group over a Galois field*, Trans. Amer. math. Soc. **71** (1951), 274-282.

[Su] S. Sundaram, Applications of the Hopf trace formula to computing homology representations, in *Jerusalem combinatorics '93. Papers from the International Conference on Combinatorics held in Jerusalem, May 9-17, 1993* (H. Barcelo and G. Kalai, editors), Contemporary Mathematics **178**, American Mathematical Society, Providence, RI, 1994.

[Su1] S. Sundaram, *The homology representations of the symmetric group on Cohen-Macaulay subposets of the partition lattice*, Adv. Math. **104** (1994), 225–296.

[Su2] S. Sundaram, *Homotopy of non-modular partitions and the Whitehouse module*, J. Algebraic. Combin., **9, No. 3** (1999), 251–269.

[Su3] S. Sundaram, *On the topology of two partition posets with forbidden block sizes*, J. Pure and Appl. Alg., **155, No. 2-3** (2001), 271–304.

[Su4] S. Sundaram, *A homotopy equivalence for partition posets related to liftings of S_{n-1} modules to S_n*, Electron. J. Combin. **9, No. 1** (2002).

[SW] S. Sundaram and M. Wachs, *The homology representations of the k-equal partition lattice*, Trans. Amer. Math. Soc., **349, No. 3**, (1997), 935–954.

[Wac] M. Wachs, *On the (co)-homology of the partition lattice and the free Lie algebra*, Discrete Math., **193** (1998), 287–319.

[Wal] J. Walker, *Topology and combinatorics of ordered sets*, Ph.D. Thesis, Massachusetts Institute of Technology, 1981.

[We] V. Welker, *Direct sum decompositions of matroids and exponential structures*, J. Combin. Theory, Ser. B **63** (1995), 222–244.

[WZZ] V. Welker, G. Ziegler, R. Zivalevic, *Homotopy types of subspace arrangements via diagrams of spaces*, Math. Ann., **295** (1993), 527–548.