

TOPOLOGICAL REPRESENTATIONS OF MATROIDS

E. SWARTZ

1. INTRODUCTION

One of the foundations of oriented matroid theory is the topological representation theorem of Folkman and Lawrence [8]. It says that an oriented (simple) matroid can be realized uniquely as an arrangement of pseudospheres. That there is no similar interpretation for the class of all matroids has been taken for granted. For instance, “A non-coordinatizable matroid of abstract origin may be thought of as a geometric object only in a purely formal way, whereas an oriented matroid may always be thought of as a geometric-topological configuration on the d -sphere (or in projective space)” [3, p. 19]. Our main theorem is that the class of geometric lattices, which is cryptomorphic to the category of simple matroids, is the same as the class of intersection lattices of arrangements of homotopy spheres.

The interpretation of a geometric lattice as an arrangement of homotopy spheres is a natural generalization of the Folkman-Lawrence theorem. An oriented matroid realizable over \mathbb{R} has a representation with geodesic spheres. Allowing pseudospheres, i.e., those which are homeomorphic, but possibly not isometric to the unit sphere, leads to the category of (simple) oriented matroids. If we further relax the conditions on the spheres to only homotopy equivalence to the standard sphere, then we are led to the category of all (simple) matroids.

Some of the theory of oriented matroids which only depends on the underlying matroid can be extended to homotopy sphere arrangements. Zaslavsky’s enumerative theory for pseudosphere arrangements can be extended to the homotopy setting. As in the oriented matroid representation theorem, the arrangement can be forced to be antipodal, so a realization as homotopy *projective* spaces is also possible. The minimal cellular resolutions of orientable matroid ideals developed in [10] can be extended to arbitrary matroids by using arrangements of homotopy spheres.

Our point of view is primarily through the lens of oriented matroids. Hence the homotopy spheres which represent the atoms of the geometric lattice have codimension one. In the future we hope to examine the point of view of complex hyperplane arrangements and consider spheres of even codimension.

The matroid theory we require is in section 2, while section 3 presents matroid Steiner complexes. We review a very general theory of arrangements of topological

Received by the editors August 29, 2002 and, in revised form, November 4, 2002.

2000 *Mathematics Subject Classification*. Primary 05B35; Secondary 52C40, 13D02, 13F55.

Key words and phrases. Matroid, geometric lattice, homotopy sphere, minimal cellular resolution.

Partially supported by a VIGRE postdoc under NSF grant number 9983660 to Cornell University.

subspaces due to Ziegler and Živaljević in section 4. Arrangements of homotopy spheres and the representation theorems are in sections 5 and 6 respectively. The last section extends the work of [10] on minimal resolutions of face ideals of independence complexes of matroids to matroid Steiner ideals.

Unlike some authors our homotopy spheres will not in general be manifolds. For a subset V of the vertices of a CW-complex Γ , the subcomplex of Γ induced by V consists of the cells of Γ all of whose incident vertices are contained in V . The rest of our notation from topology is standard. The join of two spaces X and Y is $X * Y$, their one-point union, or wedge sum, is $X \vee Y$, while $X \simeq Y$ denotes homotopy equivalence. There are three facts from topology that we will use repeatedly without specific reference. They are immediate consequences of well-known theorems of Hurewicz and Whitehead. See, for instance, [12] for proofs.

- If Γ is acyclic and simply connected, then Γ is contractible.
- If Γ is simply connected and $\tilde{H}_0(\Gamma) = \cdots = H_{i-1}(\Gamma) = \{0\}$, then $H_i(\Gamma)$ is isomorphic to $\pi_i(\Gamma)$.
- If Γ and Γ' are homotopy spheres and $f : \Gamma \rightarrow \Gamma'$ is a continuous map such that $f_* : H_*(\Gamma) \rightarrow H_*(\Gamma')$ is an isomorphism, then f is a homotopy equivalence.

2. MATROIDS

In this section we give the basic definitions and results from matroid theory that we will require. Matroid definitions and notation are as in [11]. Geometric semilattices are covered in [14]. The characteristic polynomial and Möbius invariant can be found in [18]. The beta invariant was introduced by Crapo [6].

There are numerous cryptomorphic definitions of matroids. For us a *matroid* M is a pair (E, \mathcal{I}) , E a non-empty finite ground set and \mathcal{I} a distinguished set of subsets of E . The members of \mathcal{I} are called the *independent* subsets of M and are required to satisfy:

- The empty set is in \mathcal{I} .
- If B is an independent set and $A \subseteq B$, then A is an independent set.
- If A and B are independent sets such that $|A| < |B|$, then there exists an element $x \in B - A$ such that $A \cup \{x\}$ is independent.

Matroid theory was introduced by Whitney [15]. The prototypical example of a matroid is a finite subset of a vector space over a field \mathbf{k} with the canonical independent sets. Another source of matroids is graph theory. The cycle matroid of a graph is the matroid whose finite set is the edge set of the graph and whose independent sets are the acyclic subsets of edges. Most matroid terminology can be traced back to these two types of examples.

The *circuits* of a matroid are its minimal dependent sets. If every circuit has cardinality at least three, then the matroid is *simple*. A maximal independent set is called a *basis*, and a subset which contains a basis is a *spanning* subset. The maximal non-spanning subsets are the *hyperplanes* of M . Every basis of M has the same cardinality. The *rank* of M , or $r(M)$, is that common cardinality. The *deletion* of M at e is denoted $M - e$. It is the matroid whose ground set is $E - \{e\}$ and whose independent sets are simply those members of \mathcal{I} which do not contain e . The *contraction* of M at e is denoted M/e . It is a matroid whose ground set is also $E - \{e\}$. If e is a dependent element of M , then $M/e = M - e$. Otherwise, a subset I of $E - \{e\}$ is independent in M/e if and only if $I \cup \{e\}$ is independent in M .

Deletion and contraction for a subset A of E is defined by repeatedly deleting or contracting each element of A . The restriction of M to A is $M - (E - A)$, which we shorten to $M|A$ or frequently just A . The rank of A is $r(M|A)$, which we shorten to $r(A)$. Note that $r(\emptyset) = 0$.

The dual of M is M^* . It is the matroid whose ground set is the same as M and whose bases are the complements of the bases of M . For example, $U_{r,n}$ is the matroid defined by $E = \{1, 2, \dots, n\}$ and $\mathcal{I} = \{A \subseteq E : |A| \leq r\}$. So, $U_{r,n}^* = U_{n-r,n}$. A circuit of M^* is a cocircuit of M . The complement of a cocircuit is a hyperplane.

The free extension of M is $F(M)$. It is the matroid with ground set $\tilde{E}(M) = E(M) \cup \{\tilde{e}\}$, where $\tilde{e} \notin E(M)$, and whose independent subsets are all $A \subseteq \tilde{E}(M)$ such that $|A| \leq r(M)$ and $A - \{\tilde{e}\}$ is independent in M . The free coextension of M is $(F(M^*))^*$.

A subset of M is closed if adding any element to the subset increases its rank. The closed subsets of M are also called flats. Examples of flats include the hyperplanes, also called the coatoms of M . The closed subsets of M with their inherited rank function form a ranked partially ordered set under inclusion, which we denote by $L(M)$. Given two flats X and Y in $L(M)$ their meet is $X \wedge Y = X \cap Y$. The meet of two flats is also their greatest lower bound in $L(M)$. The least upper bound of X and Y is their join, $X \vee Y$, and is equal to $X \cup Y \cup \{e : \text{there is a circuit } C, e \in C, C - \{e\} \subseteq X \cup Y\}$. When the elements of a poset are topological spaces we rely on context to clarify whether $X \vee Y$ is their poset join or one-point union.

Definition 2.1. A finite ranked poset L is a **geometric lattice** if

- (a) L is a lattice i.e., every pair of elements have a greatest lower bound and least upper bound. In particular L has a least element $\hat{0}$ and greatest element $\hat{1}$.
- (b) Every element of L other than $\hat{0}$ is the join of atoms of L .
- (c) The rank function r is semimodular, $r(X) + r(Y) \geq r(X \wedge Y) + r(X \vee Y)$.

Proposition 2.2 ([11]). *If M is matroid, then $L(M)$ is a geometric lattice. Conversely, suppose L is a geometric lattice. Let E be the atoms of L . For any $A \subseteq E$ let $\bigvee A$ be the (poset) join of all the atoms in A . Then $\mathcal{I} = \{A \subseteq E : r(\bigvee A) = |A|\}$ are the independent subsets of a matroid M such that $L = L(M)$.*

The lower interval $[\hat{0}, X] \subseteq L(M)$ is the geometric lattice of $M|X$. An upper interval $[X, \hat{1}]$ is isomorphic to the geometric lattice of M/X . For this reason $[X, \hat{1}]$ is frequently denoted L/X .

A pointed geometric lattice is a pair (L, e) with e a specified atom of L . An isomorphism of pointed geometric lattices $\phi : (L, e) \rightarrow (L', e')$ is a lattice isomorphism such that $\phi(e) = e'$. When $L = L(M)$ we will write $L(M, e)$ for the geometric semilattice $\{X \in L : e \not\leq X\}$. (Geometric semilattices are called generalized affine matroids in [9].) Even with the addition of a maximum element, $L(M, e)$ is usually not a geometric lattice.

Given any locally finite poset L the Möbius function on L is the function $\mu : L \times L \rightarrow \mathbb{Z}$ which satisfies:

- (a) If $X \not\leq Y$, then $\mu(X, Y) = 0$.
- (b) For any X in L , $\mu(X, X) = 1$.
- (c) For any X, Y in L , $X < Y$, $\sum_{X \leq Z \leq Y} \mu(X, Z) = 0$.

The *characteristic polynomial* of a geometric lattice L is

$$p(L; t) = \sum_{X \in L} \mu(\hat{0}, X) t^{r(\hat{1})-r(X)}.$$

The *beta invariant* of a geometric lattice is

$$\beta(L) = (-1)^r \sum_{X \in L} r(X) \mu(\hat{0}, X).$$

3. MATROID STEINER COMPLEXES

Matroid Steiner complexes were introduced in [5]. We follow the presentation in [4]. Let M be a matroid and $e \in E$. The *port* of M at e is the set

$$\mathcal{P} = \{C - \{e\} : C \text{ is a circuit of } M \text{ which contains } e\}.$$

The matroid Steiner complex of (M, e) is

$$\mathfrak{S}(M, e) = \{F \subseteq E - \{e\} : \forall P \in \mathcal{P}, P \not\subseteq F\}.$$

Independence complexes of matroids are a subclass of matroid Steiner complexes. If M is a matroid, then the independence complex of M is

$$\Delta(M) = \{F \subseteq E : F \text{ is independent}\}.$$

Given a matroid M let \tilde{M} be the free coextension of M with \tilde{e} the extra point. Then $\mathfrak{S}(\tilde{M}, \tilde{e}) = \Delta(M)$.

Theorem 3.1 ([4]). *Let $e \in E$, M a rank- r matroid. Then*

$$\mathfrak{S}(M, e) \simeq \bigvee_{i=1}^{\beta(M)} S^{r-2}.$$

Note that if $\beta(M)$ is zero, then $\mathfrak{S}(M, e)$ is contractible.

4. ARRANGEMENTS OF SUBSPACES

The theory of arrangements of subspaces of a topological space as presented in [19] will play a large part in our theory. Here we present a CW-version of this theory.

Let $\mathcal{A} = \{A_1, \dots, A_m\}$ be a set of distinct subcomplexes of a finite CW-complex Γ . Assume that \mathcal{A} is closed under intersections. Let P be the poset (\mathcal{A}, \leq) where $A_i \leq A_j$ if and only if $A_j \subseteq A_i$. For any $p \in P$, $\Delta(P_{<p})$ is the order complex of the poset $P_{<p} = \{q \in P : q < p\}$.

Theorem 4.1 ([19]). *Let $\mathcal{V} = \bigcup_{i=1}^m A_i$. Suppose that for every $A_j < A_i$, the inclusion map $A_i \hookrightarrow A_j$ is null-homotopic. Also assume that if A_j has multiple components, then all non-empty images $A_i \hookrightarrow A_j$ are contained in the same component. Then*

$$(1) \quad \mathcal{V} \simeq \bigvee_{i=1}^m (\Delta(P_{<A_i}) * A_i).$$

Note that $X * \emptyset = X$.

5. ARRANGEMENTS OF HOMOTOPY SPHERES

A *homotopy d-sphere* is a d -dimensional CW-complex which is homotopy equivalent to S^d . It is convenient to let the empty set be a homotopy (-1) -sphere.

Definition 5.1. A *d-arrangement of homotopy spheres* consists of a d -dimensional homotopy sphere S and a finite set of subcomplexes $\mathcal{A} = \{S_1, \dots, S_n\}$ of S each of which is a homotopy $(d - 1)$ -sphere. In addition,

- (a) Every intersection of homotopy spheres in \mathcal{A} is a homotopy sphere.
- (b) If X is an intersection in \mathcal{A} which is a d' -dimensional homotopy sphere and $X \not\subseteq S_j$, then $X \cap S_j$ is a $(d' - 1)$ -dimensional homotopy sphere.

Arrangements of homotopy spheres are a natural generalization of the arrangements of pseudospheres associated with oriented matroids. Many notions from pseudosphere arrangements can be carried over to the homotopy sphere case. The *link* of \mathcal{A} is $\mathcal{V} = \bigcup_{j=1}^n S_j$. If X is a non-empty intersection in \mathcal{A} , the *contraction* of \mathcal{A} to X is \mathcal{A}/X . It is the $(\dim X)$ -arrangement of homotopy spheres defined by letting X be the ambient homotopy sphere and setting \mathcal{A}/X equal to the collection of intersections $X \cap S_j, X \not\subseteq S_j$. The *deletion* $\mathcal{A} - S_j$ is the arrangement $\{S_1, \dots, \hat{S}_j, \dots, S_n\}$. We call \mathcal{A} *essential* if $\bigcap_{j=1}^n S_j = \emptyset$. As in the oriented matroid case, the intersection lattice of \mathcal{A} plays a key role.

Definition 5.2. Let \mathcal{A} be an arrangement of homotopy spheres. The **intersection lattice** of \mathcal{A} is the poset $L(\mathcal{A})$ of intersections of elements of \mathcal{A} ordered by reverse inclusion. By convention, $S \in L(\mathcal{A})$ as the empty intersection.

As usual, $L(\mathcal{A})/S_j \cong L(\mathcal{A}/S_j)$. We omit the elementary proof of the following.

Proposition 5.3. *Let \mathcal{A} be a d -arrangement of homotopy spheres. Then $L(\mathcal{A})$ is a geometric lattice with rank function $r(X) = d - \dim X$. If \mathcal{A} is essential, then the rank of $L(\mathcal{A})$ is $d + 1$.*

In section 6 we will prove the converse: every rank- $(d + 1)$ geometric lattice is isomorphic to $L(\mathcal{A})$ for some essential d -arrangement of homotopy spheres. As with their pseudosphere counterparts, the homotopy type of the link of \mathcal{A} only depends on $L(\mathcal{A})$.

Proposition 5.4. *If \mathcal{A} is a d -arrangement of homotopy spheres, then*

$$\mathcal{V} \simeq \bigvee_{i=1}^{|p(L(\mathcal{A});-1)|-1} S^{d-1}.$$

Proof. We apply 4.1 to all of the subcomplexes in $L(\mathcal{A})$ other than S . Let $X \in L(\mathcal{A}), X \neq S$ and let $r(X)$ be the rank of X in $L(\mathcal{A})$. Then $\Delta(L(\mathcal{A})_{Y < X})$ is the order complex of the geometric lattice $[S, X] \subseteq L(\mathcal{A})$. Hence [7],

$$\Delta(L(\mathcal{A})_{Y < X}) \simeq \bigvee_{i=1}^{|\mu(S,X)|} S^{r(X)-2}.$$

Since $X \simeq S^{d-r(X)}$, Theorem 4.1 implies

$$\mathcal{V} \simeq \bigvee_{X>S} \bigvee_{i=1}^{|\mu(S,X)|} S^{d-r(X)} * S^{r(X)-2} = \bigvee_{i=1}^{|p(L(\mathcal{A});-1)|-1} S^{d-1}.$$

□

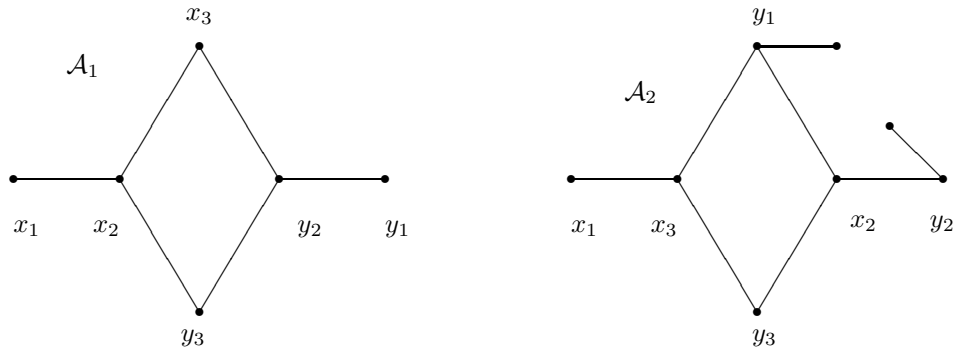


FIGURE 1. Two representations of the three-point line where $S_j = \{x_j, y_j\}$.

In contrast to pseudosphere arrangements, an arrangement of homotopy spheres may not have a “natural” cell structure. However, if S does have a CW-structure which is related to \mathcal{A} , then many of the enumerative invariants which describe the cell decomposition of a pseudosphere arrangement still hold in this more general situation.

Example 5.5. In Figure 1 both \mathcal{A}_1 and \mathcal{A}_2 are essential 1-arrangements of homotopy spheres with intersection lattice isomorphic to the three-point line. Only \mathcal{A}_1 has the same number of cells in each dimension as an essential pseudosphere arrangement with the same intersection lattice.

Definition 5.6. An arrangement of homotopy spheres is **partitioned** if the $(d-1)$ -skeleton of S is contained in \mathcal{V} . If every contraction of \mathcal{A} is partitioned, then \mathcal{A} is **fully partitioned**.

The CW-structure induced by a pseudosphere arrangement is always fully partitioned. Neither the contraction nor the deletion of a partitioned arrangement need be partitioned. For instance, \mathcal{A}_2 in Figure 1 is the deletion of a partitioned homotopy sphere arrangement with the intersection lattice of the 4-point line.

Proposition 5.7. *If \mathcal{A} is a partitioned d -arrangement of homotopy spheres, then S has $|p(L(\mathcal{A}); -1)|$ d -dimensional cells.*

Proof. The proof is virtually identical to the corresponding statements in [16]. For each cell c in S let $\psi(c) = \max\{X \in L(\mathcal{A}) : c \subseteq X\}$. For $X \in L(\mathcal{A})$ define

$$\tau(X) = \sum_{c \in \psi^{-1}(X)} (-1)^{\dim c}.$$

Note that we consider the empty set to be a cell of dimension minus one. Then, since each X is a homotopy $(d - r(X))$ -sphere,

$$\sum_{Y \geq X} \tau(Y) = \chi(X) - 1 = (-1)^{d-r(X)}.$$

Möbius inversion implies that

$$\tau(X) = \sum_{Y \geq X} \mu(X, Y)(-1)^{d-r(X)}.$$

Since \mathcal{A} is partitioned, $\tau(S)$ is $(-1)^d$ times the number of d -dimensional cells. \square

Corollary 5.8. *If every contraction $\mathcal{A}/X, r(X) \leq i$, is partitioned, then the number of $(d - i)$ -dimensional cells is*

$$\sum_{\substack{r(X)=i \\ X \leq Y}} |\mu(X, Y)|.$$

Proof. Under these conditions each $(d - i)$ -cell is in exactly one rank- i flat and $\tau(X)$ is $(-1)^{d-r(X)}$ times the number of $(d - i)$ -cells in X . \square

Unlike pseudosphere arrangements, $S - S_j$ need not consist of two contractible components. Yet, under certain conditions, it is possible to recover Zaslavsky’s enumerative results for the complex of bounded cells in a pseudosphere arrangement.

Definition 5.9. An essential arrangement of homotopy spheres \mathcal{A} is **regular** with respect to $S_j \in \mathcal{A}$ if:

- For every $X \in L(\mathcal{A})$ with $X \not\subseteq S_j$, the subcomplex of X induced by the vertices in $X - S_j$ consists of two contractible components.
- For every coatom $Y > X$ such that $X \not\subseteq S_j$ the two vertices of Y are in different components of $X - S_j$.

When \mathcal{A} is regular with respect to S_j a *bounded subcomplex* of (\mathcal{A}, S_j) is one of the two components of the subcomplex induced by the vertices in $S - S_j$. In this situation we define

$$L(\mathcal{A}, S_j) = \{X \in L(\mathcal{A}) : X \not\subseteq S_j\}.$$

Proposition 5.10. *Suppose \mathcal{A} is a partitioned d -arrangement of homotopy spheres which is regular with respect to S_j . Then the number of d -cells of a bounded subcomplex of (\mathcal{A}, S_j) is*

$$(2) \quad \left| \sum_{X \in L(\mathcal{A}, S_j)} \mu(S, X) \right| = \beta(L(\mathcal{A})).$$

Proof. Let B be a bounded subcomplex of (\mathcal{A}, S_j) . For each X in $L(\mathcal{A}, S_j)$ let $X_B = B \cap X$. For each cell $c \subseteq B$, define $\psi(c)$ to be the maximal $X \in L(\mathcal{A}, S_j)$ such that $c \subseteq X_B$. Define

$$\tau(X) = \sum_{\substack{c \in \psi^{-1}(X) \\ c \neq \emptyset}} (-1)^{\dim c}.$$

Since each X_B is contractible, its Euler characteristic is one. Therefore, for any $X \in L(\mathcal{A}, S_j)$

$$\sum_{\substack{Y \geq X \\ Y \in L(\mathcal{A}, S_j)}} \tau(Y) = 1.$$

Möbius inversion and the fact that $\tau(S)$ is $(-1)^d$ times the number of d -cells in B imply that the number of d -cells is the left-hand side of (2). That this equals $\beta(L(\mathcal{A}))$ is [16, p. 77]. \square

Corollary 5.11. *If \mathcal{A} is a fully partitioned d -arrangement of homotopy spheres which is regular with respect to S_j , then the number of $(d - i)$ -cells in a bounded subcomplex of S is*

$$\left| \sum_{\substack{Y \in L(\mathcal{A}, S_j) \\ r(X)=i}} \mu(X, Y) \right|.$$

As is evident from the proofs, the enumerative results in this section only depend on the Euler characteristic of spheres and contractible spaces. Hence these results would apply to arrangements of any spaces with the same Euler characteristics. This idea is explored in much greater generality in [17].

Not every invariant which only depends on the underlying matroid carries over from pseudosphere arrangements to homotopy sphere arrangements. The flag f -vector of a pseudosphere arrangement only depends on the intersection lattice [2]. However, as the following example shows, this is not true for homotopy sphere arrangements even if the arrangement is fully partitioned and a regular CW-complex.

Example 5.12. Let \mathcal{A} be the 2-arrangement obtained by intersecting the unit sphere in \mathbb{R}^3 with the xy, yz and xz coordinate hyperplanes. The resulting cell structure on S^2 is combinatorially equivalent to the octahedron. The intersection lattice of this arrangement is the Boolean algebra with three atoms. There are six 0-cells, twelve 1-cells, eight 2-cells, and twenty-four 0-cell \subset 2-cell incidences. Removing any single triangle and replacing it with a square along one of the coordinate hyperplanes results in a 2-arrangement of homotopy spheres with the same intersection lattice and cell counts, but twenty-five 0-cell \subset 2-cell incidences.

6. THE REPRESENTATION THEOREM

Theorem 6.1. *Let (L, e) be a rank- r pointed geometric lattice. There exists a fully partitioned essential $(r - 1)$ -arrangement of homotopy spheres \mathcal{A} which is regular with respect to S_1 such that (L, e) and $(L(\mathcal{A}), S_1)$ are isomorphic as pointed lattices. Furthermore, the arrangement can be constructed so that there exists a fixed-point free involution of S which preserves \mathcal{A} .*

Proof. The proof is by induction on r . Let $L = L(M)$. It will be evident from the construction that (\mathcal{A}, S_1) has the following additional property. Let a_1, \dots, a_m be the coatoms of L , and let $\{x_1, y_1\}, \dots, \{x_m, y_m\}$ be the corresponding zero-spheres in \mathcal{A} . Then the map which takes $X \in L(M, e)$ to the subcomplex induced by the vertices $\{x_i : X \leq a_i\}$ is a lattice isomorphism between $L(M, e)$ and $L(\mathcal{A}, S_1)$.

We begin with $r = 2$. While simply putting pairs of antipodal points around a circle satisfies the theorem, we prefer to give a procedure which produces all possible arrangements which satisfy the theorem as it is indicative of what happens in higher ranks. Since L is rank two, it consists of $\hat{0}, \hat{1}$ and coatoms (= atoms) $\{a_1, \dots, a_m\}$, where $a_1 = e$. So S has $2m$ vertices, which we label $x_1, y_1, \dots, x_m, y_m$, where the pair $\{x_i, y_i\}$ corresponds to the atom a_i . Choose any spanning tree D_b of $\{x_2, \dots, x_m\}$ and extend it to a spanning tree D of $\{x_1, \dots, x_m, y_1\}$. Let D' be the mirror image of D with $\{y_1, \dots, y_m, x_1\}$ replacing $\{x_1, \dots, x_m, y_1\}$. Finally, let S be $D \cup D'$. Now, S deformation retracts to the circle formed by the unique path from x_1 to y_1 in D concatenated with the unique path from y_1 to x_1 in D' . In addition, $\mathcal{A} = \{\{x_1, y_1\}, \dots, \{x_m, y_m\}\}$ is a 1-arrangement of homotopy zero-spheres which satisfies the theorem, where the involution is the one induced by switching x_i and




FIGURE 2. A 1-arrangement of homotopy spheres representing the 5-point line.

y_i . We also note that any arrangement which satisfies the theorem must be of this form. Figure 2 shows all three stages of this construction for the five-point line.

Now suppose the rank of L is three. Let a_1, \dots, a_m be the coatoms of L and let e_1, \dots, e_n be the atoms of L with $e_1 = e$. As above, we choose vertices $x_1, y_1, \dots, x_m, y_m$ with the pair $\{x_i, y_i\}$ corresponding to a_i . The interval $[e_1, \hat{1}]$ in L is a rank-2 geometric lattice. Therefore, we can use the construction above to obtain an arrangement of zero spheres on a homotopy 1-sphere S_1 which satisfies the theorem and represents $[e_1, \hat{1}]$ with vertices consisting of all the x_i and y_i such that $e_1 < a_i$.

Let e_j be another atom of L and let

$$W_j = \{x_i : e_j < a_i\} \cup \{y_i : a_i = e_1 \vee e_j\}.$$

As above we choose a spanning tree V_j on W_j which is also a spanning tree on the vertices of W_j which do not represent $e_1 \vee e_j$. Let \mathcal{V} be the union of S_1 and all the V_j as e_j runs through the atoms of L other than e_1 . Let \mathcal{V}_b be the subcomplex of \mathcal{V} induced by all vertices representing $\{a_i : e_1 \not\leq a_i\}$. Since \mathcal{V}_b is a connected graph, we can attach 2-cells so that the resulting space, which we call D_b , is contractible. The CW-complex $\mathcal{V} \cup D_b$ is homotopy equivalent to a connected graph, so we can attach 2-cells to it so that the resulting space is contractible. Call this two-dimensional CW-complex D . Note that none of the 2-cells attached to $\mathcal{V} \cup D_b$ have their boundary completely contained in D_b . Otherwise D would not be acyclic. Hence the subcomplex of D induced by $\{x_i : x_i \notin S_1\}$ is \mathcal{V}_b . Let D' be the mirror image of D obtained by switching the roles of x_i and y_i and let S be the union of D and D' glued along S_1 . The intersection $D \cap D' = S_1$, so S is a homotopy 2-sphere. For each $i, 2 \leq i \leq n$, let S_i be the union of V_i and its mirror image in D' . By construction each S_i is a homotopy one-sphere and the arrangement $\mathcal{A} = \{S_1, \dots, S_n\}$ satisfies the theorem.

Example 6.2. Figure 3 shows one possible way of constructing \mathcal{V} for $(L, 1)$ where L is the Fano plane as pictured. The 1-cells which will be used to form S_1, S_2 and S_3 are also labeled. The subcomplex induced by $\{D, E, F, G\}$ is \mathcal{V}_b and is the 1-skeleton of a tetrahedron. Attaching triangles EFG, DEF and DFG is one of infinitely many ways to construct $\mathcal{V} \cup D_b$.

For the induction step, assume the rank of L is greater than three. As above, let a_1, \dots, a_m be the coatoms of L and let e_1, \dots, e_n be the atoms of L with $e_1 = e$. In addition, let $x_1, y_1, \dots, x_m, y_m$ be vertices with the pair $\{x_i, y_i\}$ corresponding to a_i . The interval $L_1 = [e_1, \hat{1}]$ in L is a rank- $(r - 1)$ geometric lattice. Using the inductive algorithm we construct an arrangement \mathcal{A}_1 of homotopy $(r - 2)$ -spheres which satisfy the theorem for (L_1, e') , where e' is any atom of L_1 .

Now consider the corank-2 flats of L . Those which contain e_1 are already represented in \mathcal{A}_1 . For any other corank-2 flat X , we construct a spanning tree V_X on

$$W_X = \{x_i : X < a_i\} \cup \{y_i : a_i = X \vee e_1\},$$

which is also a spanning tree on the vertices of W_X which do not represent $X \vee e_1$.

Proceeding inductively on the corank, for each corank- k flat X neither equal to $\hat{0}$ nor above e_1 we construct a contractible space V_X on the vertices $W_X = \{x_i : X < a_i\} \cup \{y_i : X \vee e_1 \leq a_i\}$ by adding only $(k - 1)$ -dimensional cells to the complex associated to flats of lower corank. In addition, we make sure that the subcomplex of V_X induced by the vertices $\{x_i : e_1 \not\leq a_i\}$ is also contractible. Flats which lie above e_1 are already represented in S_1 .

Let \mathcal{V} be the union of S_1 and all the V_X . Let \mathcal{V}_b be the subcomplex of \mathcal{V} induced by the vertices whose corresponding coatoms do not lie above e_1 . Consider the covering of \mathcal{V}_b by the subcomplexes $U_j = V_{e_j} \cap \mathcal{V}_b, 2 \leq j$. By construction, all non-void intersections of the members of the cover are contractible so the classical nerve theorem applies. Since a collection of the U_j have void intersection if and only if the poset join of the corresponding atoms contains e_1, \mathcal{V}_b is homotopy equivalent to

$$\Delta(L, e) = \{F \subseteq \{e_2, \dots, e_m\} : e_1 \notin \bigvee F\}.$$




FIGURE 3. The Fano plane.

This complex is just the matroid Steiner complex $\mathcal{S}(M, e)$, where M is the matroid associated to L as described in Proposition 2.2. Hence, by Theorem 3.1 it has the homotopy type of a wedge of $\beta(L)$ spheres of dimension $(r - 2)$. Now we attach $\beta(L)$ cells of dimension $r - 1$ to \mathcal{V}_b in any way which kills $H_{r-2}(\mathcal{V}_b)$ and call this space D_b . Since D_b is acyclic and simply connected, it is contractible.

Now consider the cover of \mathcal{V} with the subcomplexes V_{e_j} and S_1 . Apply Theorem 4.1 to the arrangement of subcomplexes consisting of this cover and all of its intersections. Call this arrangement \mathcal{B} . For each $B \in \mathcal{B}$ define $\phi(B)$ to be the flat in L which is the intersection of all the coatoms whose corresponding vertices are incident to B . Then ϕ is a lattice isomorphism from \mathcal{B} to L . The elements of \mathcal{B} are of

two types. If $e_1 \notin \phi(B)$, then B is contractible. Otherwise, if $e_1 \leq \phi(B)$, then B is a homotopy $r - r(\phi(B))$ -sphere. As $\Delta(\mathcal{B}_{<B}) \simeq \Delta(L_{<\phi(B)})$, every non-contractible term of (1) is a wedge of $(r - 2)$ -spheres.

Since D_b is contractible and $D_b \hookrightarrow D_b \cup \mathcal{V}$ is a cofibration, the reduced homology of $D_b \cup \mathcal{V}$ is the same as the homology of the pair $(\mathcal{V}, \mathcal{V}_b)$. As \mathcal{V} is an $(r - 2)$ -dimensional CW-complex, $H_{r-2}(\mathcal{V}, \mathcal{V}_b)$ is torsion-free. This plus the long exact sequence of the pair implies that $(\mathcal{V}, \mathcal{V}_b)$, and hence $D_b \cup \mathcal{V}$, has the homology of a wedge of $(r - 2)$ -spheres. Let D be any $(r - 1)$ -dimensional CW-complex obtained by gluing $(r - 1)$ -dimensional cells to $D_b \cup \mathcal{V}$ so that the resulting space is acyclic. As before, none of these cells have their boundary contained in D_b . Hence the subcomplex induced by $\{x_i : x_i \notin S_1\}$ is \mathcal{V}_b . Since \mathcal{V} is simply connected, D is simply connected, and hence contractible. Let D' be the mirror image of D induced by switching x_i with y_i . Finally, let $S = D \cup D'$ glued along S_1 . Since S is the union of two contractible spaces whose intersection is a homotopy $(r - 2)$ -sphere, S is a homotopy $(r - 1)$ -sphere. Similar reasoning shows that for each flat $X \in L$ the subcomplex of S associated to X is a homotopy $(r - r(X))$ -sphere and the intersection lattice of the arrangement of homotopy spheres determined by the atoms of L is isomorphic to L . The involution induced by switching x_i and y_i for each i is fixed-point free and preserves the arrangement. The construction insures that this arrangement is fully partitioned and is also regular with respect to S_1 . \square

Remark 6.3. The involution means that we can also develop a theory of homotopy projective space arrangements.

Definition 6.4. Let $\mathcal{A} = \{S_1, \dots, S_n\}$ and $\mathcal{B} = \{T_1, \dots, T_n\}$ be d -arrangements of homotopy spheres. Then we say \mathcal{A} and \mathcal{B} are **homotopy equivalent arrangements** if there exists a lattice isomorphism $\phi : L(\mathcal{A}) \rightarrow L(\mathcal{B})$ and homotopy equivalences $f : S \rightarrow T, g : T \rightarrow S$ such that:

- (1) For each $X \in L(\mathcal{A})$, $f(X) \subseteq \phi(X)$, and $f : X \rightarrow \phi(X)$ is a homotopy equivalence.
- (2) For each $Y \in L(\mathcal{B})$, $g(Y) \subseteq \phi^{-1}(Y)$ and $g : Y \rightarrow \phi^{-1}(Y)$ is a homotopy equivalence.

Theorem 6.5 (Uniqueness). *Let \mathcal{A}, \mathcal{B} be essential d -arrangements of homotopy spheres such that $L(\mathcal{A}) \cong L(\mathcal{B})$. Then \mathcal{A} and \mathcal{B} are homotopy equivalent arrangements.*

Proof. Let $\phi : L(\mathcal{A}) \rightarrow L(\mathcal{B})$ be a lattice isomorphism. For every cell of S define $\psi(c)$ to be $\max\{X \in L(\mathcal{A}) : c \subseteq X\}$. By symmetry it is sufficient to construct f in the above definition. We will build up f by defining maps f_i on the i -skeletons of S inductively which satisfy the following properties:

- The restriction of f_{i+1} to the i -skeleton is f_i .
- For each corank- $(i + 1)$ flat $X \in L(\mathcal{A})$, $f_i : X \rightarrow \phi(X)$ is a homotopy equivalence.
- If $\dim c \leq i$, then $f_i(c) \subseteq \phi(\psi(c))$.

Since the arrangements are essential, the coatoms are all homeomorphic to two disjoint points. Choose a homeomorphism h of the union of the coatoms of \mathcal{A} to the coatoms of \mathcal{B} which preserves ϕ . Now extend h to f_0 by arbitrarily choosing any image point in $\phi(\psi(v))$ for any vertex v which is not in a coatom of \mathcal{A} .

Now assume that f_{i-1} has been defined. Let X be a corank- $(i + 1)$ flat of $L(\mathcal{A})$. Since \mathcal{A} is essential, X is a homotopy i -sphere. Let c be an i -cell of X . The definition of f_{i-1} insures that $f_{i-1}(\partial(c)) \subseteq \phi(X)$. As $\phi(X)$ is also a homotopy i -sphere, there is a map $f_c : \bar{c} \rightarrow \phi(X)$ such that f_c equals f_{i-1} when restricted to $\partial(c)$. Putting all of these maps together gives a map $f_X : X \rightarrow \phi(X)$. The induced map in homology, $(f_X)_\star : H_i(X) \rightarrow H_i(\phi(X))$, is multiplication by n_X after choosing generators for the respective homology groups. If $n_X = \pm 1$, then f_X is a homotopy equivalence. If not, choose any i -cell c in X and redefine f_c as follows. Let $\alpha : (D^i, S^{i-1}) \rightarrow X$ be the attaching map for c . Let $D^i(1/2)$ be the closed ball of radius one-half and let $S^{i-1}(1/2) = \partial D^i(1/2)$. Replace f_c with \tilde{f}_c which satisfies:

- \tilde{f}_c restricted to $S^{i-1}(1/2)$ is constant and equal to $f_c(0)$. The induced map $(\tilde{f}_c)_\star : H_i(D^i(1/2), S^{i-1}(1/2)) \rightarrow H_i(\phi(X))$ is multiplication by $1 - n_X$ with respect to the appropriate generators.
- $\tilde{f}_c(\mathbf{x}) = f_c((2|\mathbf{x}| - 1)\mathbf{x})$ for $|\mathbf{x}| \geq 1/2$.

The new f_X induces an isomorphism on homology and hence is a homotopy equivalence. If an i -cell c is not contained in any corank- $(i + 1)$ flat, then $\psi(c)$ is at least i connected, so we define an arbitrary map $f_c : c \rightarrow \phi(\psi(c))$ which equals f_{i-1} on ∂c . Putting all of the f_c together gives the required map f_i . □

7. MINIMAL CELLULAR RESOLUTIONS OF MATROID STEINER IDEALS

One approach to finding syzygies of monomial ideals is through minimal cellular resolutions. The following presentation of minimal cellular resolutions is taken from [10]. Let \mathbf{k} be a field and let I be the monomial ideal $\langle m_1, \dots, m_s \rangle$ in the polynomial ring $\mathbf{k}[x_1, \dots, x_n]$, which we denote by $\mathbf{k}[\mathbf{x}]$. Let Γ be a CW-complex with s vertices v_1, \dots, v_s which are labeled with the monomials m_1, \dots, m_s . Write $c \geq c'$ whenever a cell c' belongs to the closure of another cell c . Label each cell c of Γ with the monomial $m_c = \text{lcm}\{m_i : v_i \leq c\}$, the least common multiple of the monomials labeling the vertices of c . Set $m_\emptyset = 1$ for the empty cell of Γ . The principal ideal $\langle m_c \rangle$ is identified with the free \mathbb{N}^n -graded $\mathbf{k}[\mathbf{x}]$ -module of rank 1 with generator in degree $\text{deg } m_c$. For a pair of cells $c \geq c'$, let $p_c^{c'} : \langle m_c \rangle \hookrightarrow \langle m_{c'} \rangle$ be the inclusion map of ideals. It is a degree-preserving homomorphism of \mathbb{N}^n -graded modules.

Fix an orientation of each cell in Γ , and define the *cellular complex* $C_\star(\Gamma, I)$

$$\dots \xrightarrow{\partial_3} C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} C_{-1} = \mathbf{k}[\mathbf{x}]$$

as follows. The \mathbb{N}^n -graded $\mathbf{k}[\mathbf{x}]$ -module of i -chains is

$$C_i = \bigoplus_{c: \dim c=i} \langle m_c \rangle,$$

where the direct sum is over all i -dimensional cells c of Γ . The differential $\partial_i : C_i \rightarrow C_{i-1}$ is defined on the component $\langle m_c \rangle$ as the weighted sum of the maps $p_c^{c'}$:

$$\partial_i = \sum_{\substack{c' \leq c \\ \dim c' = i-1}} [c : c'] p_c^{c'},$$

where $[c : c'] \in \mathbb{Z}$ is the incidence coefficient of the oriented cells c and c' in the usual topological sense. The differential ∂_i preserves the \mathbb{N}^n -grading of $\mathbf{k}[\mathbf{x}]$ -modules. Note that if $m_1 = \dots = m_s = 1$, then $C_\star(\Gamma, I)$ is the usual chain complex of Γ over $\mathbf{k}[\mathbf{x}]$. For any monomial $m \in \mathbf{k}[\mathbf{x}]$, define $\Gamma_{\leq m}$ to be the subcomplex

of Γ consisting of all cells c whose label m_c divides m . We call any such $\Gamma_{\leq m}$ an I -essential subcomplex of Γ .

Proposition 7.1 ([1, Proposition 1.2]). *The cellular complex $C_\star(\Gamma, I)$ is exact if and only if every I -essential subcomplex of Γ is acyclic over \mathbf{k} . Moreover, if $m_c \neq m_{c'}$ for any $c > c'$, then $C_\star(\Gamma, I)$ gives a minimal free resolution of I .*

If both conditions of the above proposition are met, then we call Γ an I -complex and $C_\star(\Gamma, I)$ a *minimal cellular resolution* of I . Recall that $\beta_i(I)$ is the \mathbf{k} -dimension of the i^{th} free module in a minimal free resolution of I . When Γ is an I -complex the number of i -dimensional cells in Γ is $\beta_i(I)$.

Given an abstract simplicial complex Δ with vertices v_1, \dots, v_n the *face ideal* of Δ in $\mathbf{k}[\mathbf{x}]$ is

$$I_\Delta = \langle \{x_{i_1} \cdots x_{i_s} : \{v_{i_1}, \dots, v_{i_s}\} \notin \Delta \} \rangle.$$

When Δ is a matroid Steiner complex we call I_Δ a *matroid Steiner ideal*. As pointed out in section 3.1, independence complexes of matroids are a special subclass of matroid Steiner complexes. The problem of finding minimal resolutions of I_Δ when Δ is the independence complex of a matroid M was examined in [10]. When M is an orientable matroid Novik et al. showed that the bounded subcomplex of any pseudosphere arrangement which realizes M^\star extended by a free point is an I_Δ -complex.

Let I be a matroid Steiner ideal. The topological representation theorem allows a complete description of all possible equivalence classes of complexes which are I -complexes for every field \mathbf{k} . Two I -complexes are *equivalent* if they have the same cellular resolution (up to orientation). Acyclic 2-complexes which are not simply connected show that it is possible for two equivalent I -complexes to be homotopy inequivalent.

Theorem 7.2. *Let $I = I_{S(M, e)}$ be a matroid Steiner ideal. Let (\mathcal{A}, S_1) be a fully partitioned arrangement of homotopy spheres which is regular with respect to S_1 such that $(L(\mathcal{A}), S_1) \cong (L(M^\star), e)$. Then a bounded subcomplex of (\mathcal{A}, S_1) is an I -complex. Conversely, if Γ is an I -complex over every field k , then Γ is such a complex.*

Proof. Let Γ be a bounded subcomplex of $L(\mathcal{A}, S_1)$ and let ϕ be a pointed lattice isomorphism from $(L(\mathcal{A}), S_1)$ to $(L(M^\star), e)$. For notational simplicity we assume that $E = \{1, \dots, n\} = [n]$ and $e = 1$. As usual, for each cell c of Γ let $\psi(c) = \max\{X \in \mathcal{A} : c \subseteq X\}$. Since \mathcal{A} is essential, $\psi(c)$ is the meet of the coatoms of $L(\mathcal{A})$ which correspond to the vertices of c . If v is a vertex in Γ , then $\phi(v)$ is a coatom of $L(M^\star)$ which does not contain e . Similarly, for each cell c of Γ , $\phi(\psi(c))$ is a flat of $L(M^\star)$ which does not contain e . Label each cell c with m_c , the square-free monomial whose support is $([n] - \phi(\psi(c))) - \{1\}$. Matroid duality implies that the support of m_c is the union of the circuits of M which are the complements of the coatoms corresponding to the vertices incident to c . Thus each cell of Γ is labeled with $\text{lcm}\{m_i : v_i \leq c\}$. As the rank of any flat X of $L(\mathcal{A})$ is equal to $d - \dim(X)$, Γ satisfies $m_c \neq m_{c'}$ whenever $c > c'$. Applying matroid duality again, we see that every I -essential subcomplex is of the form $\phi^{-1}(X) \cap \Gamma$, where $X \in L(M^\star, e)$. The regularity of \mathcal{A} with respect to S_1 guarantees that every I -essential subcomplex of Γ is contractible, and hence acyclic.

For the converse, assume that Γ is an I -complex over every field k . Then each I -essential subcomplex of Γ is acyclic over \mathbb{Z} . Relabel each cell c with the complement in $[n] - \{1\}$ of the support of m_c , i.e., the flat of M^* which does not contain e and corresponds by matroid duality to $\text{lcm}_{v \leq c} m_v$. The zero-skeleton of Γ is the same as the zero-skeleton of any bounded subcomplex of (\mathcal{A}, S_1) , where \mathcal{A} is an essential arrangement of homotopy spheres which is regular with respect to S_1 and $(L(\mathcal{A}), S_1) \cong (L(M^*), e)$. As noted above, all I -essential subcomplexes of Γ consist of cells whose labels contain a fixed $X \in L(M^*, e)$. Proceeding inductively on the corank of all the flats in $L(M^*, e)$ we see that Γ must be equivalent to one constructed in exactly the same fashion as the procedure in the representation theorem for constructing D_b for $L(M^*, e)$. So, D_b and the simultaneously constructed (\mathcal{A}, S_1) are the required complexes. \square

Corollary 7.3. *Let $I = I_{S(M,e)}$ be a matroid Steiner ideal. Then*

$$\beta_i(I) = \left| \sum_{\substack{r(X)=n-r-i \\ e \notin Y}} \mu_{L(M^*)}(X, Y) \right|$$

Proof. Theorem 7.2 and Corollary 5.11. \square

When $\Delta = \Delta(M)$ is the independence complex of a matroid, Corollary 7.3 recovers Stanley’s formula [13]

$$\beta_i(I_\Delta) = \left| \sum_{r(X)=n-r-i} \mu_{L(M^*)}(X, \hat{1}) \right|.$$

In this case $\Delta = \mathfrak{S}(\tilde{M}, \tilde{e})$, where \tilde{M} is the free coextension of M and \tilde{e} is the extra point. The poset of flats of \tilde{M}^* which do not lie above \tilde{e} is $L(M^*)$ with $\hat{1}$ removed. So,

$$\mu_{L(M^*)}(X, \hat{1}) = - \sum_{\substack{Y \in L(\tilde{M}^*) \\ \tilde{e} \notin Y}} \mu_{L(\tilde{M}^*)}(X, Y).$$

The method developed by Novik et al. to construct minimal resolutions for face ideals of independence complexes of matroids can be applied to matroid Steiner ideals. Instead of applying [10, Corollary 3.10] to the order dual of $L(M^*)$, we can use the order dual of $\hat{L}(M^*, e) = L(M^*, e) \cup \{\hat{1}\}$. The only non-trivial point is that every upper interval of $\hat{L}(M^*, e)$ is Cohen-Macaulay. This follows from the fact that for any $X \in \hat{L}(M^*, e)$ the upper interval $[X, \hat{1}]$ is isomorphic to $\hat{L}(M^*/X, e)$ and every geometric semilattice is shellable and hence Cohen-Macaulay [14].

Equivalence classes of CW-complexes which are I -complexes over every field can also be parameterized algebraically. Following the notation of [10, p. 299], view the complex $\mathcal{Z}(P)$ as a complex over \mathbb{Z} . Use P equal to the order dual of $\hat{L}(M^*, e)$. By Theorem 7.2 equivalence classes of such I -complexes come from all the possible D_b constructed in the representation theorem for $(L(M^*), e)$. Every D_b corresponds to choosing bases for $H_*(\mathcal{V}_X)$ for each $X \in L(M^*, e)$. Working backwards from Z_{-1} and using the fact that ϕ is an injection when restricted to each direct sum component, this is equivalent to choosing a \mathbb{Z} -basis for every direct sum component which occurs in $\mathcal{Z}(P)$. The only restriction to these bases is that the image under ϕ of any basis of a component in $\bigoplus_{rk(F)=2} Z_0(\Delta(F))$ must be the difference of

exactly two basis elements of $\bigoplus_{rk(F)=1} Z_{-1}(\Delta(F))$. This is a reflection of the fact that the boundary of any one-cell is always the difference of two zero-cells.

ACKNOWLEDGEMENTS

Bernd Sturmfels suggested the problem of determining all I -complexes when I is the face ideal of the independence complex of a matroid. Louis Billera, Tom Zaslavsky and the anonymous referee helped clear up several points in the exposition. Tom Zaslavsky also pointed out the relevance of [14].

REFERENCES

- [1] D. Bayer and B. Sturmfels. Cellular resolutions of monomial ideals. *J. Reine Angew. Math.*, 502:123–140, 1998. MR **99g**:13018
- [2] M. Bayer and B. Sturmfels. Lawrence polytopes. *Canadian J. Math.*, 42:62–79, 1990. MR **91e**:52023
- [3] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. Ziegler. *Oriented matroids*. Cambridge University Press, second edition, 1999. MR **2000j**:52016
- [4] M.K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Mathematics*, 217:101–113, 2000. MR **2001g**:52106
- [5] C.J. Colbourn and W.R. Pulleybank. Matroid Steiner problems, the Tutte polynomial and network reliability. *J. Comb. Theory Ser. B*, 41:20–31, 1989. MR **90g**:05059
- [6] H. Crapo. A higher invariant for matroids. *J. Combinatorial Theory*, 2:406–417, 1967. MR **35**:6579
- [7] J. Folkman. The homology groups of a lattice. *J. Math. Mech.*, 15:631–636, 1966. MR **32**:5557
- [8] J. Folkman and J. Lawrence. Oriented matroids. *J. Comb. Theory Ser. B*, 25:199–236, 1978. MR **81g**:05045
- [9] J. McNulty. Generalized affine matroids. In *Proceedings of the twenty-fifth Southeastern international conference on combinatorics, graph theory and computing*, volume 101 of *Congressus Numerantium*, pages 243–254, 1994. MR **96k**:05034
- [10] I. Novik, A. Postnikov, and B. Sturmfels. Syzygies of oriented matroids. *Duke Math. J.*, 111(2):287–317, 2002.
- [11] J. G. Oxley. *Matroid Theory*. Oxford University Press, Oxford, 1992. MR **94d**:05033
- [12] E.H. Spanier. *Algebraic Topology*. McGraw-Hill, 1966. MR **35**:1007
- [13] R.P. Stanley. Cohen-Macaulay complexes. In M. Aigner, editor, *Higher combinatorics*, pages 51–62, 1977. MR **58**:28010
- [14] M. Wachs and J. Walker. On geometric semilattices. *Order*, 2(4):367–385, 1986. MR **87f**:06004
- [15] H. Whitney. On the abstract properties of linear dependence. *American Journal of Mathematics*, 57:509–533, 1935.
- [16] T. Zaslavsky. Facing up to arrangements: face-count formulas for partitions of space by hyperplanes. *Mem. Amer. Math. Soc.*, 1(1):154, 1975. MR **50**:9603
- [17] T. Zaslavsky. A combinatorial analysis of topological dissections. *Adv. in Math.*, 25:267–285, 1977. MR **56**:5310
- [18] T. Zaslavsky. The Möbius function and the characteristic polynomial. In N.L. White, editor, *Combinatorial geometries*. Cambridge University Press, 1987.
- [19] G. Ziegler and R. Živaljević. Homotopy types of subspace arrangements via diagrams of spaces. *Math. Ann.*, 295(3):527–548, 1993. MR **94c**:55018

MALOTT HALL, CORNELL UNIVERSITY, ITHACA, NEW YORK 14853
E-mail address: ebs@math.cornell.edu