AN INTRODUCTION TO THE MOEBIUS FUNCTION

C. D. Godsil [1]

Combinatorics and Optimization
University of Waterloo
Waterloo, Ontario
Canada N2L 3G1

[1] Support from grant OGP0009439 of the National Sciences and Engineering Council of Canada is gratefully acknowledged.

ABSTRACT

This is an introduction to the Möbius function of a poset. The chief novelty is in the exposition. We show how order-preserving maps from one poset to another can be used to relate their Möbius functions. We derive the basic results on the Möbius function, applying them in particular to geometric lattices.

C. D. Godsil

1. Posets and Matrices

Our first four sections provide a fairly standard approach to the Möbius function of a poset. It is based in part on the treatment in Chapter 2 of Lovász [12].

Let P be a poset with elements $p_1, \ldots, p_n$. (Unless we explicitly say otherwise, all posets we consider are finite. So n is an honest-to-God Kroneckerian integer.) An $n \times n$ matrix B is *compatible* with P if $(B)_{ij}$ is zero unless $p_i \le p_j$. It is immediate that set of all $n \times n$ matrices compatible with P is closed under addition, and it is not hard to show that it is also closed under multiplication. Thus it is an algebra over $\mathbf{C}$, often called the *incidence algebra* of P. We note that it contains the identity matrix, as well as the *Zeta matrix* Z_P, which has ij-entry equal to one if and only if $p_i \le p_j$. Any matrix compatible with P can be regarded as a function on $P \times P$. This function is non-zero only on ordered pairs (x, y) where $x \le y$, and so we may even view our function as a function on the *intervals* of P.

A simple induction argument shows that, by relabelling he elements of P if needed, we may assume that $i \le j$ whenever $p_i \le p_j$. Then the matrices compatible with B are all upper triangular, and so such a matrix is invertible if and only if its diagonal entries are all non-zero. We have the following interesting result.

1.1 Lemma. *Let P be a poset. If B is compatible with P and invertible then B^{-1} is compatible with P.*

Proof. Let $\varphi(x)$ be the characteristic polynomial of B. If B is invertible then $\varphi(0) \ne 0$ and so

$$\varphi(x) = x\psi(x) + c$$

for some polynomial ψ and non-zero constant c. By the Cayley-Hamilton theorem

$$0 = \varphi(B) = B\psi(B) + cI,$$

whence $c^{-1}\psi(B) = B^{-1}$. □

Since the zeta matrix Z_P has all its diagonal entries equal to one, it is invertible. By the lemma, $(Z_P)^{-1}$ is compatible with P. The corresponding function on $P \times P$ is the *Möbius function* of P, and is denoted by μ_P.

We can determine μ_P by inverting the triangular matrix Z_P; this represents no intellectual challenge and can be carried out in polynomial time. However, for many interesting posets, properties of the Möbius function can be read from properties of the poset. The values taken by the Möbius function may have combinatorial significance.

2. Möbius Inversion

Our first result is known as the principle of Möbius inversion.

2.1 Theorem. *Let P be a poset and let f and g be functions on P. Then*
 a) $g(x) = \sum_{y\ge x} f(y)$ *if and only if* $f(z) = \sum_y \mu(z,y)g(y)$.
 b) $g(x) = \sum_{y\le x} f(y)$ *if and only if* $f(z) = \sum_y \mu(y,z)g(y)$.

Proof. We may abuse notation and view f and g as column vectors, with entries indexed by P. Then (a) says that

$$g = Z_P f \Leftrightarrow M_P g = f$$

and (b) that

$$g = Z_P^T f \Leftrightarrow M_P^T g = f.$$

Since $M_P = Z_P^{-1}$, no more need be said. □

Since all diagonal entries of Z_P are equal to one, it follows that the same is true for M_P. (One way to convince yourself of this is to recall that the diagonal entries of a triangular matrix are its eigenvalues, and that the eigenvalues of Z_P^{-1} are the reciprocals of the eigenvalues of Z_P.) Thus $\mu_P(x,x) = 1$, for any element x of P. There is a recursive expression for the remaining values of μ_P, equivalent to the back-substitution phase in Gaussian elimination.

2.2 Lemma. *Let a and b be two elements of the poset P. Then*

$$\mu_P(a,b) = \begin{cases} 0, & \text{if } a \not\le b; \\ 1, & \text{if } a = b; \\ -\sum_{x:x\le b}\mu_P(a,x), & \text{otherwise.} \end{cases}$$

Proof. If $a \not\le b$ then $(M_P)_{ab} = 0$, since M_P is compatible with P. If $a = b$ then $\mu_P(a,a) = 1$, as noted above. Finally, if $a < b$ then $(Z_P M_P)_{ab} = 0$ and therefore

$$0 = \sum_{x\le b} \mu_P(a,x).$$

Hence $\mu(a,b) = -\sum_{z<b}\mu_P(a,x)$, as required. □

The argument used in the previous proof yields another useful identity. Suppose a and b are elements of the poset P and $a < b$. Then

$$\mu_P(a,b) = -\sum_{x:a<x} \mu_P(x,b).$$

The *chain* $\mathcal{C}(n)$ is the poset with elements $0,\ldots,n$, where $i \geq j$ if $i-j$ is non-negative. Suppose a and b are elements of $\mathcal{C}(n)$ and $\mu = \mu_{\mathcal{C}(n)}$. If $a < b$ then $\mu(a,b) = -1$ if b covers a, and is zero otherwise. We will use this in the next section to compute the Möbius function for the poset of divisors of a given integer.

3. Products

The *product* of posets P and Q is the poset with elements $P \times Q$, where

$$(x,y) \leq_{P\times Q} (x',y')$$

if and only if

$$x \leq_P x' \text{ and } y \leq_Q y'.$$

We consider two examples. Let $\mathcal{B}(n)$ be the lattice of subsets of an n-element set. It is routine to verify that $\mathcal{B}(n)$ is isomorphic to the product of n copies of $\mathcal{B}(1)$, which in turn is isomorphic to $\mathcal{C}(1)$. The lattice of divisors of an integer n is also isomorphic to a product of chains. More precisely, if p is prime and $n = p^r$ then the lattice of divisors of n is the chain of length r. If $n = p^r m$ where m and p are coprime then the divisor lattice of n is the product of the divisor lattice of m with the chain of length r. Note that $\mathcal{B}(n)$ can be regarded as the divisor lattice of a square-free integer having exactly n distinct prime divisors.

Turning from examples to Möbius functions, we have

$$Z_{P\times Q} = Z_P \otimes Z_Q,$$

whence

$$M_{P\times Q} = M_P \otimes M_Q.$$

As an immediate consequence we have the next result.

3.1 Lemma. *If P and Q are posets and (x,y) and (x',y') are elements of $P \times Q$ then*

$$\mu_{P\times Q}((x,y),(x',y')) = \mu_P(x,x')\mu_Q(y,y').$$

Suppose that S and T are subsets of some n-element set. Then, taken with our remarks above, the previous lemma implies that

$$\mu(S,T) = \begin{cases} 0, & S \not\subseteq T; \\ (-1)^{|T\setminus S|}, & \text{otherwise.} \end{cases}$$

We now present a classical combinatorial application of the Möbius function. A *derangement* is a permutation with no fixed points. We wish to compute D_n, the number of derangements of n points.

To this end, if $S \subseteq \{1,\ldots,n\}$ let $D_n(S)$ denote the number of permutations of $\{1,\ldots,n\}$ which fix each point in S and no points not in S. (So $D_n(\emptyset) = D_n$.) Let $F_n(S)$ denote the number of permutations which fix each point in S. Both $F_n(S)$ and $D_n(S)$ are functions on $\mathcal{B}(n)$. We have

$$F_n(S) = (n - |S|)!$$

and we will use this to compute D_n.

The key observation is that

$$F_n(S) = \sum_{T\supseteq S} D_n(T)$$

whence

$$\begin{aligned} D_n(S) &= \sum_T \mu_{\mathcal{B}(n)}(S,T)\, F_n(T) \\ &= \sum_{T\supseteq S} (-1)^{|T\setminus S|}\, F_n(T) \end{aligned}$$

Assuming that $|S| = k$, we may write the last sum as

$$\sum_{\ell=k}^{n} (n-\ell)! \binom{n-k}{\ell-k} (-1)^{\ell-k}.$$

and therefore

$$\begin{aligned} D_n &= \sum_{\leq 0}^{n} (n-\ell)! \binom{n}{\ell} (-1)^{\ell} \\ &= n! \left(1 - \frac{1}{1!} + \frac{1}{2!} - \cdots + (-1)^n \frac{1}{n!}\right) \\ &= \left[\frac{n!}{e}\right] \end{aligned}$$

4. Posets and Chains

A *chain* in a poset is a set of elements, any two of which are comparable. Any finite chain has unique minimal and maximal elements. The set of all non-empty chains of the poset P will be denoted by $\mathrm{Ch}(P)$. This set is partially ordered by inclusion, hence is itself a poset. Our first task in this section is to describe the relation between chains and the Möbius function. We denote the length of the chain C by $\ell(C)$. (This is one less than the number of elements of C.)

For this we need another definition. If P is a poset with elements $p_1, \dots, p_n$, let Y_P be the $n \times n$ matrix with ij-entry equal to one if and only if $p_i < p_j$. Thus, if we have arranged things so that Z_P is triangular then $Y_P = Z_P - I$.

4.1 Lemma. *Let P be a poset with elements $p_1, \dots, p_n$. Then*
(a) the ij-entry of Y_P^m is the number of chains of length m in P with least element p_i and maximal element p_j,
(b) the ij-entry of Z_P^m is a polynomial in m, and
(c) the ij-entry of M_P is $\sum_{C \in \mathrm{Ch}(P)} (-1)^{\ell(P)}$.

Proof. Given that $Y_P^{m+1} = Y_P^m Y$, it is easy to prove (a) by induction on m. If $m > |P|$ then $Y_p^m = 0$. Assuming that $Z_P = I + Y_P$, we then have

$$Z_P^m = \sum_{k=0}^{m} \binom{m}{k} Y_P^k.$$

Since $Y_P^k = 0$ for sufficiently large k and since $\binom{m}{k}$ is a polynomial in m (of degree k), it follows that the entries of Z_P^m are polynomials in m.

To prove (c) we observe that

$$M_P = Z_P^{-1} = (I + Y_P)^{-1} = \sum_{k \geq 0} (-1)^k Y_P^k. \qquad \square$$

Lemma 4.1(c) is quite important, and it is worth recording it in a slightly different form.

4.2 Lemma (P. Hall). *If a and b are elements of the poset P then*

$$\mu_P(a, b) = \sum (-1)^{\ell(C)},$$

where the sum is over all chains C in P with minimal element a and maximal element b. $\square$

We can always create a new poset from P by reversing the order. The result is a poset, P^{op} say, with the same elements as P such that

$$a \leq_P b \Leftrightarrow b \leq_{P^{\mathrm{op}}} a.$$

One immediate consequence of Hall's theorem is that

$$\mu_P(a,b) = \mu_{P^{\mathrm{op}}}(b,a).$$

This can often be used to derive alternate forms of various identities, e.g., the expression for μ_P we gave directly following Lemma 2.2 can be derived from Lemma 2.2 in this way.

Making use of terminology to be explained later, the ab-entry of $(Z_P)^m$ can be shown to be equal to the number of order preserving mappings from a chain of length m into P. (The corresponding entry of $(Y_P)^m$ counts order preserving injections.)

5. Simplicial Complexes

A simplicial complex $\mathcal{S}$ on a set Ω is a set of non-empty subsets of Ω such that if $A \in \mathcal{S}$ and $B \subseteq A$ then $B \in \mathcal{S}$. (Oh well, there are two schools of thought. Some authors choose to make the empty set an element of any simplicial complex.) The elements of S are called *faces* and the *dimension* of a face A is $|A| - 1$. (Yes, the empty set would have dimension -1.) The maximal elements of $\mathcal{S}$ are called *simplices*. We denote the number of k-dimensional faces of $\mathcal{S}$ by $f_k(\mathcal{S})$, and call it the *k-th level number of the complex*. The *Euler characteristic* of $\mathcal{S}$ is defined to be

$$\sum_{k\geq 0}(-1)^k f_k.$$

We consider two examples. Let $\mathcal{M}$ be a triangulation of a surface and let $\mathcal{S}$ be the simplicial complex whose elements are the sets of vertices contained in some face of $\mathcal{M}$. (To be more prosaic, the elements of $\mathcal{S}$ are the vertices, edges and triangles of $\mathcal{M}$.) In this case the Euler characteristic of $\mathcal{S}$ is determined by the surface on which $\mathcal{M}$ lies.

Our second example is $\mathrm{Ch}(P)$. The simplices are the maximal chains in P. If $\widehat{P}$ is obtained from P by adjoining a new 0- and 1-element then the Euler characteristic of $\mathrm{Ch}(P)$ is equal to

$$1 + \mu_{\widehat{P}}(0,1).$$

To make matters more confusing, we note that every simplicial complex is a poset. We will see later that $\mathcal{S}$ and $\mathrm{Ch}(\mathcal{S})$ have the same Euler characteristic.

6. Determinants

The theory we describe in this section is one of the prettiest parts of the theory of the Möbius function, and was developed independently by Lindström [11] and Wilf [18].

6.1 Lemma. *Let f be a function defined on the poset P and set*

$$g(x,y) = \sum_{z \geq x,y} f(z).$$

If G is the matrix with rows and columns indexed by P and xy-entry equal to $g(x,y)$ then $\det G = \prod_{x \in P} f(x)$.

Proof. Let F be the diagonal matrix with rows and columns indexed by P, where $(F)_{xx} = f(x)$. Then $G = Z_P F Z_P^T$ and so

$$\det G = \det(Z_P F Z_P^T) = (\det Z_P)^2 \det F = \det F. \qquad \square$$

Exercise: give an expression for f in terms of g.

If P is a lattice then $g(x,y) = \sum_{z \geq x \vee y} f(z)$. Thus we may allow g to be any function defined on P, with f given by

$$f(y) = \sum_z \mu_P(y,z) g(z).$$

Then Lemma 6.1 implies that

$$\det G = \prod_{x \in P} \sum_{y \in P} \mu_P(x,y) g(y). \tag{6.1}$$

We will make significant use of this result later. Further applications appear in the papers of Lindström and Wilf [11,18] and in [12: Chapter 2].

7. Order-Preserving Mappings

Let P and Q be posets. A function f from P to Q is *order preserving* if, whenever x and y belong to P and $x \leq y$, we have $f(x) \leq f(y)$. To consider one example, if P is $\mathcal{B}(n)$ and Q if the chain of length n then the mapping from P to Q which sends each set to its cardinality is order preserving. In this section we will see how an order preserving mapping can be used to establish a relation between μ_P and μ_Q.

To begin, we introduce the *Möbius number* of a poset. If P is a poset, let $\widehat{P}$ be the poset obtained from P by adjoining a new zero-element $\widehat{0}$ and a new one-element $\widehat{1}$. Hence if $x \in P$ then

$$\widehat{0} <_{\widehat{P}} x <_{\widehat{P}} \widehat{1}.$$

We define the Möbius number $\mu(P)$ of P by

$$\mu(P) := \mu_{\widehat{P}}(\widehat{0}, \widehat{1}).$$

It is equal to the number of chains of even length in P, less the number of chains of odd length. Note that the Möbius number of the empty poset is -1. (Why?)

The following simple result will be one of our main tools. It implies that if a poset P has a 1-element then $\mu(P) = 0$.

7.1 Lemma. *If the poset P has an element which is comparable with all elements of P then $\mu(P) = 0$.*

Proof. Suppose that a is comnparable with all elements of P. Then there is bijection between the chains in P which contain a and those which do not. □

If $a \in P$ and a is conparable with every element of P, we will say that P is a *cone* over a.

More definitions. Suppose that P is a poset and $a \in P$. By $P_{a\leq}$ we denote the set of elements x of P such that $a \leq x$, while $P_{a<}$ consists of the elements x such that $x > a$. Similarly we define $P_{\leq a}$ and $P_{<a}$. Now we can state the main result of this section. It is more or less equivalent to Theorem 5.5 in Baklawski [2].

7.2 Theorem (Baklawski). *Let P and Q be posets and let f be an order-preserving mapping from P to Q. Then*

$$\mu(Q) = \mu(P) + \sum_{y \in Q} \mu(Q_{y<})\mu(f^{-1}(Q_{\leq y})).$$

A poset of the form $f^{-1}(Q_{\leq y})$ will be called a *fibre* of f. Note that is a subset of P. The poset $Q_{\leq y}$ has a 1-element and so its Möbius number is zero (by Lemma 7.1). If all fibres of f have 1-elements then it follows from Theorem 7.2 that $\mu(P) = \mu(Q)$. A subset P of a poset S is an *ideal* if, whenever $a \in S$ and $x \leq a$, we have $x \in P$. Any fibre of an order-preserving mapping is an ideal. It is worth noting that if f is an order-preserving mapping from S to the chain $\mathrm{Ch}(1)$ then $f^{-1}(0)$ is an ideal and, conversely, each ideal of S determines an order-preserving mapping into $\mathrm{Ch}(1)$. A subset F of S is a *filter* if whenever $x \geq a$ and $a \in F$, we have $x \in F$.

The following result is a consequence of Theorem 7.2, but we give a direct proof of it.

7.3 Lemma. *If P is an ideal of the poset S then*

$$\mu(S) = \mu(P) + \sum_{y \in S \setminus P} \mu(S_{y<})\mu(P_{\leq y}). \tag{7.1}$$

Proof. We use Lemma 4.2. Suppose that C is a chain in S. If $C \subseteq P$ then, in the right side of $\widehat{\text{m}}$apa, it is counted by the term $\mu(P)$. If $C \not\subseteq P$, let y be the least element of $C \setminus P$. Then $C \setminus y$ is the disjoint union of a chain from $P_{\leq y}$ and a chain from $S_{y<}$. It is easy to check that in this case C is counted, with the correct sign, by the expression $\mu(S_{y<})\mu(P_{\leq y})$. □

Now we show how to derive Theorem 7.2 from Lemma 7.3. Assume that f is an order-preserving map from P to Q. Construct a new poset S with element set $P \cup Q$ by declaring that $a \leq b$ if either

(a) $a, b \in P$ and $a \leq_P b$, or
(b) $a, b \in Q$ and $a \leq_Q b$, or
(c) $a \in P$, $b \in Q$ and $f(a) \leq_Q b$.

(This construction is due to Baklawski [2].) It is easy to see that S is a poset and P is an ideal in it. Hence we have

$$\mu(S) = \mu(P) + \sum_{y \in S \setminus P} \mu(S_{y<})\mu(P_{\leq y}). \tag{7.2}$$

If $y \in S \setminus P = Q$ then $S_{y<} = Q_{y\leq}$ and $P_{y\leq} = f^{-1}(Q_{\leq y})$, whence we deduce that

$$\mu(S) = \mu(P) + \sum_{y \in Q} \mu(Q_{y<})\mu(f^{-1}(Q_{\leq y})). \qquad \square$$

(This is just a dual version of Lemma 7.3.)

Since P is an ideal in S it follows that Q^{op} is an ideal in S^{op}. Therefore

$$\begin{aligned}\mu(S) = \mu(S^{\mathrm{op}}) &= \mu(Q^{\mathrm{op}}) + \sum_{x \in S^{\mathrm{op}} \setminus Q^{\mathrm{op}}} \mu({S^{\mathrm{op}}}_{x<})\mu({Q^{\mathrm{op}}}_{\leq x}) \\ &= \mu(Q) + \sum_{x \in P} \mu(S_{<x})\mu(Q_{x\leq}).\end{aligned}$$

As $Q_{x\leq} = Q_{f(x)\leq}$ has a 0-element for all x in P, this implies that $\mu(S) = \mu(Q)$. Hence Theorem 7.2 follows.

Exercise: Derive Lemma 7.3 from Theorem 7.2.

In [16] Walker proves a more general result than Theorem 7.2: he allows the order-preserving mapping f to be an *ideal relation* between P and Q, i.e., an ideal in $P \times Q$. This has the advantages of being more general, and more symmetric in the roles P and Q play.

8. Retracts

A mapping $f : P \mapsto P$ is *decreasing* if $f(x) \leq x$ for all x in P. A subposet Q of S is a *retract* if there is an order-preserving and decreasing mapping f from S to Q such that $f \restriction Q$ is the identity, and then we call f a *retraction*. If is order-preserving and increasing then we also define the fixed points of f to be a retract of S. Note that if S is constructed from P and Q as in the proof of Theorem 7.2 then the mapping which sends a in Q to itself and a in P to $f(a)$ is a retraction. We saw in the proof of Theorem 7.2 that $\mu(S) = \mu(Q)$. More generally we have the following result.

8.1 Lemma. *If Q is a retract of S then $\mu(Q) = \mu(S)$.*

Proof. Let f be a retraction from S onto Q. If $x \in f^{-1}(Q_{\leq y})$ then $f(x) \leq y$. As $f(y) \leq y$ (indeed $f(y) = y$) it follows that y is a 1-element in $f^{-1}(Q_{\leq y})$, hence each fibre of f has a 1-element and therefore has Möbius number zero. By Theorem 7.2 we deduce that $\mu(Q) = \mu(S)$. □

8.2 Corollary. *If P is a poset then $\mu(P) = \mu(\mathrm{Ch}(P))$.*

Proof. Each element of P is a chain, therefore P is a subposet of $\mathrm{Ch}(P)$. Consider the map f from $\mathrm{Ch}(P)$ to P defined by setting $f(C)$ equal to $\max(C)$. Then f is order-preserving, decreasing and its restriction to P is the identity. □

A *point* in a lattice is an element which covers 0.

8.3 Corollary. *Let L be a lattice. If 1 is not a join of points then $\mu_L(0,1) = 0$.*

Proof. If $x \in L \setminus 0$, define $f(x)$ to be join of the points in L below x. Then f is order-preserving and decreasing and $f(f(x)) = f(x)$ for all x in $L \setminus 0$. If $f(1) \neq 1$ then $f(1) < 1$ and F is retract of L'. Hence $\mu(F) = \mu(L')$ and, since $f(1)$ is a 1-element in F, it follows that $\mu(L') = 0$. □

If a and b are elements of a poset P and the least upper bound of a and b is defined, we denote it by $a \vee b$.

8.4 Lemma. *Let P be a poset. If $a \in P$ and $a \vee x$ exists for all x in P then $\mu(P) = 0$.*

Proof. There are two steps. First, a is a 0-element in $P_{a\leq}$ and so $\mu(P_{a\leq}) = 0$. Second, the map $x \to x \vee a$ is order preserving and increasing, with $P_{a\leq}$ as its set of fixed points. Hence $P_{a\leq}$ is a retract of P and therefore $\mu(P) = 0$. □

8.5 Lemma (Weisner). *If L is a lattice and $a \in L \setminus 0$ then*

$$\mu_L(0,1) = -\sum_{x \vee a = 1,\ x<1} \mu_L(0,x).$$

Proof. Suppose

$$G := \{x \in L' : x \vee a < 1\}.$$

Then $a \in G$ and $a \vee x$ exists for all x in G, so $\mu(G) = 0$ by the previous lemma. Since G is an ideal in L, using Lemma 2.2 we find that

$$\mu(G) = \sum_{x \in G \cup 0} \mu_L(0,x).$$

We now have

$$\begin{aligned}\mu_L(0,1) = -\sum_{x<1} \mu_L(0,x) &= -\sum_{x\vee a=1, x<1} \mu_L(0,x) - \sum_{x \in G\cup 0} \mu_L(0,x)\\ &= -\sum_{x\vee a=1, x<1} \mu_L(0,x) - \mu(G)\end{aligned}$$

This yields the lemma. □

In d̂elete, we will need the next result. The proof is left as an easy exercise.

8.6 Lemma. *Let f be an order preserving and decreasing mapping of P into itself and let F be the set of fixed points of f. Then F is a retract of P.* □

9. Cutsets

A *cutset* in a lattice L is a set C which contains at least one element from each maximal chain. Call a non-empty subset S of C a *simplex* it it has a bound (i.e., a meet or a join) in $L\setminus\{0,1\}$. The set of all simplices in C forms a simplicial complex, which we denote by $\Gamma(L,C)$. By way of example, if L is the lattice of subspaces of a finite-dimensional vector space V then the set of all 1-dimensional vector spaces is a cutset, C say. A subset of C lies in $\Gamma(L,C)$ if and only if its join is not the entire space, i.e., if and only if it is not a spanning set in V. For any lattice L, let L' denote the poset obtained from L by deleting its 0- and 1-element. Thus

$$\mu(L') = \mu_L(0,1),$$

which provides one reason why we need L'.

9.1 Theorem. *If L is a lattice and C is a cutset then $\mu(L') = \mu(\Gamma(L,C))$.*

Proof. Let $\Gamma(L,C)$ be abbreviated to Γ. If B is a chain in L', define $f(B)$ by

$$f(B) := \{x \in C : x \cup B \in \mathrm{Ch}(L')\}.$$

(In other words, $f(B)$ is the set of all elements of C which are comparable with each element of B.) Since C is a cutset, $f(B) \neq \emptyset$ and it follows that f is an order-preserving mapping from $\mathrm{Ch}(L')$ to Γ^{op}. Hence we may prove the theorem by showing that all fibres of f have Möbius number zero.

Let S be an element of Γ^{op}, and let F denote the fibre $f^{-1}(\Gamma^{\mathrm{op}}_{\leq S})$. If a chain D of L lies in this fibre then $S \subseteq f(D)$. If $x \in D$ then x is comparable with every element of S. Hence

$$\wedge S \leq x \leq \vee S.$$

Since $S \in \Gamma(L,C)$, either $\wedge S$ or $\vee S$ lies in L'. Assume $\vee S \in L'$, and denote it by z. Then for any element x of D we have $x \leq z$, whence $D \cup z \in \mathrm{Ch}(L')$. Further every element of S is comparable with all elements of $D \cup z$, thus $D \cup z$ belongs to F.

Now z is a chain in L' and $S \subseteq f(z)$. Hence $z \in F$ and $z \cup D$ lies in F for all elements of F. By Lemma 8.4 it follows that F has Möbius number zero. A similar argument yields the same conclusion when if $\wedge S \in L'$. Hence the theorem holds. $\square$

We can manipulate this theorem to obtain a more explicit formula for $\mu(L')$.

9.2 Corollary. *Let C be a cutset in the lattice L, and let a_k be the number of k-subsets of C with join 1 and meet 0. Then $\mu_L(0,1) = \sum_{k=0}^{|C|}(-1)^k a_k$.*

Proof. If $S \in \Gamma$ then the interval $\widehat{\Gamma}_{\leq S}$ is a Boolean lattice and so

$$\mu_{\widehat{\Gamma}}(0,S) = (-1)^{|S|},$$

from we find, using Lemma 2.2, that

$$\mu(\Gamma) = \mu_{\widehat{\Gamma}}(0,1) = -\sum_{S \in \widehat{\Gamma} \setminus 1} \mu_{\widehat{\Gamma}}(0,S) = 1 + \sum_{S \in \Gamma}(-1)^{|S|}. \tag{9.1}$$

Define a_k to be the number of k-subsets S of C such that $\wedge S = 0$ and $\vee S = 1$. Then, when $1 \leq k \leq |C|$, the number of k-subsets of Γ is equal to

$$\binom{|C|}{k} - a_k$$

Now, assuming $|C| > 0$ (which is the only interesting case)

$$0 = \sum_{S \subseteq C,\ S \in \Gamma} (-1)^{|S|} + \sum_{S \subseteq C,\ S \notin \Gamma,} (-1)^{|S|}.$$

From (9.1) we see that the first sum here is equal to $\mu(\Gamma) - 1$, while the second sum equals $1 + \sum_{k \geq 0}(-1)^k a_k$. The result follows. □

We consider applications of Corollary 9.2. Suppose that L is the lattice of subspaces of a finite-dimensional vector space over some finite field, and let C be the set of all 1-dimensional subspaces. Then C is a cutset and a_k is the number of spanning subsets of C with cardinality k.

For another example, let G be a graph with vertex set V and let L be the set of all partitions of V such that each cell induces a connected subgraph of G. Then the join of any two elements of L lies in L and hence L is a lattice, but not in general a sub-lattice of the lattice of all partitions of V. Let C be the set of all partitions in L with one cell of size two and all others singletons. (So the cell of size two is an edge of G.) Then C is a cutset in L and a_k is the number of subgraphs of G with k edges and the same number of connected components as G.

Rota [14] proved Theorem 9.1 under the assumption that C was a cutset and an antichain. Walker proves an even more general result than Theorem 9.2 in [16], our proof is based on his. (Our task is slightly more complicated, in that Walker can use ideal relations where we must use order-preserving functions.)

10. Complements

The main result in this section is a slight weakening of Theorem 8.1 from Walker [16]. To begin, we derive a technical lemma.

10.1 Lemma. *Let L be a lattice and suppose $s \in L'$. Let G be the set of all elements x of L' such that $x \vee s < 1$. If $\mu(G_{\leq y}) \neq 0$ then y is a complement to s.*

Proof. Note that $s \in G$. We show that if y is not a complement to s in L then $\mu(G_{\leq y}) = 0$. If $y \in G$ then $G_{\leq y}$ has a 1-element and so its Möbius number is zero. If $y \notin G$ and $y \wedge s = 0$ then y is a complement to S. Suppose $y \notin G$ and $y \wedge s \neq 0$. If $z \in G_{\leq y}$ the

$$z \vee (y \wedge s) \vee s = z \vee s < 1.$$

Hence $z \vee (y \wedge s) \in G$ and therefore $z \in G_{\leq y}$. Thus $G_{\leq y}$ is a cone over $y \wedge s$ and so has Möbius number zero. The lemma follows. □

10.2 Theorem (Walker [16]). *Let L be a lattice, let a be an element of L' and let $a^{\perp}$ be the set of all complements of a. Then $\mu(L' \setminus a^{\perp}) = 0$.*

Proof. Let M denote $L' \setminus a^{\perp}$ and let G be the subposet of L' consisting of all elements x of M such that $x \vee a < 1$. Then G is an ideal of L' and contains a. The fibres of the inclusion mapping of G in M are the sets $G_{\leq y}$, where y in M. By the previous lemma these fibres all have Möbius number zero, whence $\mu(M) = \mu(G)$. Since $a \in G$ and $a \vee x$ exists for all x in G, we see by Lemma 8.4 that $\mu(G) = 0$. It follows that $\mu(M) = 0$, as required. □

Exercise: Let s be an element of the lattices L. Let S be a subset of L' containing $a^{\perp}$ such that if $x \in S$ then $x \vee a = 1$. Show that $\mu(L' \setminus S) = 0$.

10.3 Corollary. *If L is a lattice and $\mu_L(0, 1) \neq 0$ then L is complemented.*

Proof. If some element of L has no complement then the theorem applies, with $S = \emptyset$. □

11. Topology

There is more going on than we have yet admitted. An order-preserving map f from a poset P to a poset Q induces an order-preserving map from $\mathrm{Ch}(P)$ to $\mathrm{Ch}(Q)$. But $\mathrm{Ch}(P)$ and $\mathrm{Ch}(Q)$ are simplicial complexes and thus may be viewed as topological spaces. The map induced by f is then a continuous map.

Let X and Y be topological spaces. Two continuous functions f and g from X to Y are *homotopic* if there is a continuous function

$$\Phi : X \times [0,1] \to Y$$

such that $\Phi(x,0) = f(x)$ and $\Phi(x,1)$ for all x in X. We say two topological spaces X and Y are homotopic if there are continuous functions $f : X \to Y$ and $g : Y \to X$ such that $g \circ f$ and $f \circ g$ are homotopic to the respective identity maps on X and Y. It can be shown that homotopy is an equivalance relation on topological spaces. Any convex subset of $\mathbf{R}^n$ is homotopic to a point, while two homotopic surfaces in $\mathbf{R}^3$ are homeomorphic. We say posets P and Q are homotopic if $\mathrm{Ch}(P)$ and $\mathrm{Ch}(Q)$ are.

For our purposes, the following is important.

11.1 Lemma. *If P and Q are posets such that $\mathrm{Ch}(P)$ is homotopic to $\mathrm{Ch}(Q)$ then $\mu(P) = \mu(Q)$.*

Proof. The Möbius number of P is determined by the Euler characteristic of $\mathrm{Ch}(P)$. Homotopic simplicial complexes have the same Euler characteristic. □

A topological space is *contractible* if it is homotopic to a point. One class of contractible simplicial complexes are *cones*. A simplicial complex $\mathcal{S}$ is a cone if it contains an element v such that $v \vee x$ is defined for all elements x of $\mathcal{S}$. It is not hard to see that if the poset P is a cone then $\mathrm{Ch}(P)$ is a cone as a simplicial complex (and as a poset). We will say a poset P is contractible if $\mathrm{Ch}(P)$ is. We have the following important result.

11.2 Theorem (Quillen [13]). *Let P and Q be posets and let f be an order-preserving map from P to Q. If $f^{-1}(q)$ is contractible for any element q of Q then P and Q are homotopic.* □

Note that in all cases where we have proved that the fibres of some order-preserving map have Möbius number zero, we have actually shown that the fibres are cones and hence contractible. Thus if L is a lattice, $p \in L$ and S is the set of complements of s in L then $L' \setminus S$ is contractible.

If P and Q are posets, it makes sense to talk about two order-preserving maps f and g from P to Q as being homotopic. No combinatorial characterisation of what this means is known. However if $f(x) \leq g(x)$ for all x in P then it is easy to show that f is homotopic to g.

We may view Q^P as a poset, where $f \leq g$ for two elements f and g of Q^P if $f(x) \leq g(x)$ for all x in P. The poset Q^P is the disjoint union of a number of connected components; two maps in the same component will be homotopic. The constant map taking each element of P to a fixed element of Q is always order-preserving, so $|Q^P| \geq |Q|$.

It turns out that lying in the same component of Q^P is not a good approximation to the topological notion of homotopy, for reasons we now discuss.

Suppose that a is an element of P which covers a unique element b of P. Define a map

$$\varphi_a : P \to P \setminus a$$

by setting $\varphi_a(x)$ equal to x if $x \neq a$ and $\varphi_a(a) = b$. Call φ_a a *deletion*. Each fibre of the inclusion mapping of $P \setminus a$ into P has a 1-element. (It is not hard to see that $P \setminus a$ is a retract of P.)

Now suppose that f is an order-preserving map of P into itself and $f(x) \leq x$ for all elements x of P. Let a be an element of P that is minimal, subject to the condition that $f(a) < a$. If $b < a$ then, by our choice of a, we have $f(b) = b$. On the other hand, f is order-preserving and so $f(b) \leq f(a)$. Hence if $b < a$ then $f(b) \leq f(a)$ and we have shown that $f(a)$ is the unique element of P covered by a.

Exercise: Show that any order-preserving and decreasing map from P into itself is a composition of deletions.

It is possible that, by applying a sequence of deletions, we might able to map P onto the poset with exactly one element. In this case we say that P is *dismantlable*, and P is homotopic to a 1-element poset, i.e., it is contractible.

Exercise: Let P be a poset. The following are equivalent:
(a) P is dismantlable,
(b) P^X has exactly one component for any poset X and
(c) P^P has exactly one component.

12. Geometric Lattices

The Möbius function is particularly useful when applied to geometric lattices. This section introduces these lattices briefly. There are two parts to their definition.

A lattice is a *point-lattice* if every non-zero element can be expressed as the join of points.

A lattice is *semimodular* if, whenever a and b are elements such that if a covers $a \wedge b$ then b is covered by $a \vee b$. There are a number of equivalent definitions, which we will discuss shortly. What we have just called a semimodular lattice is more strictly an upper semimodular lattice. A lattice which is dual to a semimodular lattice is *lower semimodular*.

A lattice is *geometric* if it is a semimodular point lattice. One class of examples arises as follows. Let X be a set of points in a finite-dimensional projective space. A *flat* in X is any subset of X of the form $H \cap X$, where H is a projective subspace. The lattice of flats of X is geometric.

Exercise: Show that $\mathcal{P}(n)$, the lattice of all partitions of an n-set, is geometric.

The points of a geometric lattice may also be referred to as *atoms*. A maximal flat is called a *hyperplane*.

Exercise: Show that each element in a geometric lattice is the meet of a set of hyperplanes.

We will use the result of the next exercise several times.

Exercise: Show that any interval in a geometric lattice is geometric.

For the remainder of this section, we discuss some of the properties of semimodular lattices.

12.1 Lemma. *A lattice L is semimodular if $a \vee b$ covers both a and b whenever a and b cover $a \wedge b$.* □

A poset P is *ranked* if any two maximal chains joining the same pair of elements have the same length. (Equivalently, we may say that P satisifies the Jordan-Dedekind condition.) If P is ranked and $a \in P$ then the maximum length of a chain ending on a is the rank of a, which we denote by $r(a)$. If P is a ranked poset and b covers a then $r(b) = r(a) + 1$.

12.2 Lemma. *Let L be a lattice with rank function r. Then L is semimodular if and only if*

$$r(a \wedge b) + r(a \vee b) \le r(a) + r(b) \tag{12.1}$$

for all elements a and b of L. □

We call (12.1) the *semimodular identity*. A ranked lattice is modular if equality holds in the semimodular identity for all pairs of elements a and b. The lattice of subspaces of a vector space is modular, as are the Boolean lattices.

We will need the next result in ĝeo-moeb.

12.3 Lemma. *Any point in a geometric lattice has a complement.*

Proof. Let L be geometric and let p be a point in L. Let a be an element of L which is maximal, subject to the condition that $a \wedge p = 0$. If $a \vee p = 1$ then a is a complement of p and we are finished.

Otherwise $a \vee p < 1$ and, since 1 is join of points, it follows that there is point q of L such that $q \not\le a \vee p$. Now there are two possibilities. If $p \le a \vee q$ then

$$p \vee a \le q \vee a.$$

But p and q cover 0, hence both $a \vee p$ and $a \vee q$ cover a. This implies that $r(a \vee p) = r(a \vee q)$ and therefore $q \le a \vee p$, which contradicts our choice of q.

If $p \not\le a \vee q$ then $p \wedge (a \vee q) = 0$. Since $q \not\le a \vee p$, it follows that $a < a \vee q$, and thus we have a contradiction to our choice of a. □

For further background on geometric lattices, see [1, 5].

13. Modular Elements

An element a in a geometric lattice L is *modular* if the semimodular identity holds for all pairs (a, b), i.e.,

$$r(a \wedge b) + r(a \vee b) = r(a) + r(b)$$

for all elements b of L. Equivalently a is modular if the set of all complements of L form an antichain.

Exercise: Prove that an element a in a geometric lattice is modular if and only if its complements form an antichain.

Any point in a geometric lattice L is modular. If a is a point of L and $b \in L$ then, since a covers 0, either $a \wedge b = 0$ or $a \wedge b = a$. In the first case

$$r(x) + 1 = r(x) + r(a) \geq r(x \wedge a) + r(x \vee a) = r(x \vee a) > r(x),$$

while in the second

$$r(x) + r(a) \geq r(x \wedge a) + r(x \vee a) = r(a) + r(x).$$

In both cases we have equality in the semimodular inequality. The fact that points are modular is not always useful; it may be better to have modular elements of higher rank. We note two examples.

Let $\mathcal{B}_q(n)$ denote the lattice of subspaces of an n-dimensional vector space over a field with q elements. This is a modular lattice, and thus all its elements are modular. If $\mathcal{P}(n)$ is the lattice of all partitions of $\{1, \ldots, n\}$ then the partition with cells $\{1, \ldots, n-1\}$ and $\{n\}$ is a modular hyperplane.

13.1 Lemma. *Let a and b be elements in the geometric lattice L. If a is modular then the map $x \mapsto x \vee a$ is an isomorphism from $[a \wedge b, a]$ to $[b, a \vee b]$.*

Proof. If $x \in [a \wedge b, a]$ then the mapping $x \mapsto b \wedge x$ is order-preserving, as is the mapping $y \mapsto a \vee y$ when $y \in [a, a \vee b]$. Hence the composite map ψ defined by

$$\psi(x) = a \wedge (x \vee b)$$

is an order-preserving map from $[a \wedge b, a]$ into itself. Since $a \wedge (x \vee b) \geq x$, it is also increasing.

Suppose $c \in [a \wedge b, a]$. Since $a \wedge b = c \wedge b$, the semimodular identity implies that

$$r(c \vee b) - r(b) = r(c) - r(a \wedge b).$$

Applying the semimodular identity to the pair $(a, \vee b)$ and noting that $(a \vee c) \vee b = b \vee a$, we get

$$r(a \vee b) - r(c \vee b) \leq r(a) - r(a \wedge (c \vee b)).$$

Summing the last two inequalities yields that

$$r(a \vee b) - r(b) \leq r(a) - r(a \wedge b) - (r(a \wedge (c \vee b)) - r(c)).$$

As a is modular

$$r(a \vee b) - r(b) = r(a) - r(a \wedge b)$$

and, given the previous inequality, we deduce that

$$r(a \wedge (c \vee b)) \leq r(c).$$

However $c \leq a \wedge (c \vee b)$ and therefore we have proved that $c = a \wedge (c \vee b)$. So if a is modular then ψ is the identity mapping and the intervals $[a \wedge b, a]$ and $[b, a \vee b]$ are isomorphic. □

14. Möbius Functions and Geometric Lattices

Our first result will enable us to compute the Möbius function on intervals in $\mathcal{B}_q(n)$ and $\mathcal{P}(n)$. We need one preliminary result.

14.1 Lemma. *If C is an antichain in the poset P then*

$$\mu(P) = \mu(P \setminus C) + \sum_{x \in C} \mu(P_{a<})\mu(P_{<a}).$$

Proof. Apply Theorem 7.2 with f the inclusion mapping of $P \setminus C$ into P. The details are left as an exercise. □

This lemma is useful even when C is a single element of P.

14.2 Theorem. *Suppose a is a modular element of the geometric lattice L, not* 0 *or* 1*, and let $a^\perp$ be the set of all complements of a in L. Then*

$$\mu_L(0,1) = \mu_L(0,a) \sum_{x \in a^\perp} \mu_L(0,x).$$

Proof. Let P be L' and let $a^\perp$ be the set of all complements to a in L. Then $a^\perp$ is an antichain and so, using Lemma 14.1, we get

$$\mu(L') = \mu(L' \setminus a^\perp) + \sum_{x \in a^\perp} \mu(L'_{x<})\mu(L'_{<x}).$$

By Theorem 10.2 we have that $\mu(L' \setminus a^\perp) = 0$. By Lemma 13.1, if $x \in a^\perp$ then the intervals $[x \wedge a, a] = [0, a]$ and $[x, x \vee a] = [x, 1]$ are isomorphic, hence

$$\mu(L'_{x<}) = \mu_L(x, 1) = \mu_L(0, a).$$

As $\mu(L'_{<x}) = \mu_L(0, x)$, the theorem follows. □

14.3 Corollary. *If L is a geometric lattice and a and b are elements of L such that $a \le b$ then* $(-1)^{r(b)-r(a)}\mu_L(a,b) > 0$.

Proof. Since any interval of a geometric lattice is geometric, it suffices to assume that $a = 0$ and $b = 1$. Let p be a point in L. Then p is modular and all its complements are hyperplanes. (It has complements by Lemma 12.3.) By the theorem

$$\mu_L(0,1) = -\sum_{x \in p^\perp} \mu(0,x).$$

We may assume inductively that $\mu(0,x)$ is non-zero and has the same sign for all x in $p^\perp$, whence the result follows. □

Next we compute the Möbius function on $\mathcal{B}_q(n)$ using Theorem 14.2. Let h be a hyperplane in $\mathcal{B}_q(n)$. Then h is modular and so, if $L = \mathcal{B}_q(n)$,

$$\mu_L(0,1) = \mu_L(0,h) \sum_{p \in h^\perp} \mu_L(0,p).$$

Since h is modular, all its complements are points. Consequently $\mu_L(0,p) = -1$. The number of points in $h^\perp$ is q^{n-1} and therefore

$$\mu_L(0,1) = -q^{n-1}\mu_L(0,h).$$

As $\mu_L(0,h) = \mu_{\mathcal{B}_q(n-1)}(0,1)$, a trivial induction argument yields that

$$\mu_{\mathcal{B}_q(n)}(0,1) = (-1)^n q^{\binom{n}{2}}.$$

We can also compute the Möbius function for $\mathcal{P}(n)$. Here $h = \{\{1\},\{2,\dots,n\}\}$ is a modular hyperplane whose complements are the partitions with one non-trivial cell, of the form $\{1,i\}$. Hence

$$\mu_{\mathcal{P}(n)}(0,1) = -(n-1)\mu_{\mathcal{P}(n-1)}(0,1).$$

Once again a simple induction argument yields that

$$\mu_{\mathcal{P}(n)}(0,1) = (-1)^{n-1}(n-1)!.$$

15. Broken Circuits

The main result of this section shows that if L is a geometric lattice then $(-1)^r\mu_L(0,1)$ is not only non-negative, it counts something.

Let L be a geometric lattice and let S be the set of all points in it. Since L is a point-lattice, we can identify each element of L with the set of points below it in L. We can extend the rank function of L to a function on subsets of S by defining $r(T)$ to be $r(\vee T)$, for any subset T of S. We have

(1) $r(\emptyset) = 0$,
(2) if $p \in S$ then $r(p) = 1$,
(3) if T and U are subsets of S and $T \subseteq U$ then $r(T) \leq r(U)$ and
(4) for any pair of subsets T and U of S,

$$r(T) + r(U) \geq r(T \cup U) + r(T \cap U).$$

We define a subset T of S to be *independent* if $r(T) = |T|$, all other subsets are *dependent*. The set S, together with its collection of independent subsets, is a *matroid*. A *circuit* is a minimal dependent subset of S. A *flat* is a subset, F say, of S such that if $p \in S \setminus F$ then $r(p \cup F) > r(F)$. Thus the flats correspond precisely to the elements of L. We will not be doing any matroid theory, but we will need to refer to the circuits and independent sets of a geometric lattice.

The independent sets of a geometric lattice form a simplicial complex—every subset of an independent set is independent. We are now going to define the *broken circuit complex*, which is a subcomplex of the independent set complex. Assume L is a geometric lattice and let $\trianglelefteq$ be a total order on its points. A set of points is a *broken circuit* if it can be obtained from some circuit by deleting its least element, relative to $\trianglelefteq$. The broken circuit complex $\mathrm{Br}(L)$ has as its elements all independent sets which do not contain a broken circuit. Since a set of points which contains no broken circuit cannot contain a circuit, the elements of $\mathrm{Br}(L)$ are all independent sets. If $T \in \mathrm{Br}(L)$ then $\vee T$ contains no point less than the least element of T. (If p is a point in $\vee T$ and $p \notin T$ then there is a circuit $p \cup S$, for some subset S of T.)

15.1 Theorem (Whitney [17]). *Let L be a geometric lattice and let $\trianglelefteq$ be a total order on its points. Then the number of independent sets of k points which contain no broken circuit is $\sum_{a:r(a)=k}(-1)^k\mu_L(0,a)$.*

Proof. Any independent set of size k lies in a unique element of L with height k. Hence it suffices to prove that $(-1)^r\mu_L(0,1)$ is the number of independent sets of r atoms containing no broken circuits, where r is the height of L. We prove this by induction on r.

Let p_1 be the least point of L and let b be a complement of p in L. Then b has height $r-1$. Let M the geometric lattice formed by the interval $[0,b]$ and let $\mathrm{Br}(M)$ be the broken circuit complex of M, relative to the ordering of the points of M obtained by restriction of $\trianglelefteq$. We claim that T is an independent set of $r-1$ points of M then $T\cup p_1 \in \mathrm{Br}(L)$ if and only if $T\in \mathrm{Br}(M)$.

Suppose first that $T\in \mathrm{Br}(M)$. If $T\cup p_1$ contains a broken circuit C from L then either $C\subseteq T$ and so C is a broken circuit in M, or $C\cup p_1$ is a circuit in L and therefore

$$p_1 \in \vee C \le b.$$

Conversely, let S be an r-subset in $\mathrm{Br}(L)$. Then all points of L lie in $\vee S$, and therefore $p_1\cup S$ is dependent. Since S is independent any circuit in $p_1\cup S$ must contain p_1, whence S contains a broken circuit.

By induction, the number of $(r-1)$-subsets of $\mathrm{Br}(M)$ is equal to

$$(-1)^{r-1}\mu_M(0,1) = (-1)^{r-1}\mu_L(0,b)$$

and therefore the number of r-sets in $\mathrm{Br}(L)$ is equal to

$$(-1)^{r-1}\sum_{b\in p_1^\perp}\mu_L(0,b).$$

Since $\mu_L(0,p_1) = -1$, by Theorem 13.2 this last sum is equal to $(-1)^r\mu_L(0,1)$, as required. $\square$

Exercise: A simplicial complex is *pure* if all its maximal elements have the same height. Show that any broken circuit complex is pure.

16. The Partition Lattice

In this section we apply Theorem 15.1 to the partition lattice $\mathcal{P}(n)$. We identify the points of $\mathcal{P}(n)$ with the edge set of K_n and we let $\trianglelefteq$ denote the lexicographic order on the points. An independent set is then a forest, i.e., an acyclic subgraph of K_n. If $i < j < k$ then the edges ik and jk form a broken circuit. It follows that a forest F in $E(K_n)$ contains no broken circuit if and only if each component of F has the property that the vertices in any path going away from the least vertex form an increasing sequence. (This condition is equivalent to containing no "broken triangle", the details are up to you.) Consequently the forests in K_n containing no broken circuits can be viewed as non-increasing functions on the set $\{1, \ldots, n\}$, the number of components in the forest is equal to the number of fixed points of the function.

Since 1 is a fixed point of any non-increasing function, the number of such functions with exactly one fixed point is $(n-1)!$. This shows that

$$\mu_{\mathcal{P}(n)}(0, 1) = (-1)^{n-1}(n-1)!.$$

Fortunately this is consistent with our earlier result. The problem which remains is to determine the number of non-increasing functions with exactly k fixed points when $k > 1$. I claim that this is equal to the number of permutations of $\{1, \ldots, n\}$ with exactly k cycles.

The proof of this is indirect. The first step is an encoding of a permutation in cyclic form as a sequence. Start with a permutation in cyclic form, with any 1-element cycles written out explicitly. (To give an extreme case, the identity permutation in cyclic form is usually written as (1), but we must write it as $(1)(2)\cdots(n)$.) Now write each cycle so that the largest element is first (so if $n = 4$ then (123) is now (312)(4)). Next, order the cycles so that the first elements form an increasing sequence. Finally remove the parentheses. You are invited to prove that we have now defined a bijection from $\mathrm{Sym}(n)$ onto itself. Denote the image of a permutation β under this bijection by $\widehat{\beta}$.

What is the relation between the cycles of β and $\widehat{\beta}$? If $\sigma \in \mathrm{Sym}(n)$, define j to be a *record* if $\sigma(i) < \sigma(j)$ whenever $i < j$. Our claim (well it is my claim, but you have to prove it) is that the number of cycles in β is equal to the number of records of $\widehat{\beta}$.

But this only completes the first step; we need to convert $\widehat{\beta}$ into a non-decreasing function. This is easy. If $\sigma \in \mathrm{Sym}(n)$, let f_σ be defined by

$$f_\sigma(j) = |\{i : i < \sigma^{-1}(j),\ \sigma(i) \leq j\}|$$

Again, you must convince yourself that σ can be reconstructed from f_σ. Note however that the fixed points of f_σ are precisely the records of σ. Hence the number of fixed points of $f_{\widehat{\beta}}$ is equal to the number of cycles of β.

A forest in K_n with exactly $n-d$ edges has exactly d components. So we have shown that the number of forests with k edges which contain no broken circuit is equal to the number of permutations of $\{1,\ldots,n\}$ with exactly $n-k$ cycles. This number has no nice explicit form, but it is known to be equal to $(-1)^k$ times the coefficient of x^k in $x(x-1)\cdots(x-n+1)$. The coefficient itself is a *Stirling number of the first kind.* (For background see, e.g., [15: Chapter 1].)

17. Contractions and Colourings

We consider a family of geometric lattices including $\mathcal{P}(n)$ as a special case. Let G be a graph with vertex set V and edge set E. A *contraction* of G will be defined to be a partition of V such that the subgraph induced by any cell is connected. Equivalently we may view them as subsets S of E with the property that, for any edge $f \in E \setminus S$, the number of components of $S \cup f$ is less than the number of components of S. We will be a question and denote the set of all contractions of G by L_G. Every contraction of G is a partition of V and so, if $n = |V|$, it can viewed as an element of $\mathcal{P}(n)$. Further the join of any two contractions is a contraction, and so the contractions of G form a sub-semilattice of $\mathcal{P}(n)$. As we remark in the Appendix, any join-semilattice with zero can be turned into a lattice—in this case we define the meet of contractions σ and τ by

$$\sigma \wedge \tau := \vee\{\gamma \in L_G : \gamma \leq \sigma, \tau\}.$$

Exercise: Show that L_G, as defined above is a geometric lattice, and that if S is a set of points of L_G then $n - r(S)$ is the number of components of S.

The points of L_G are precisely the edges of G. The independent sets of points are precisely the (edge-sets of the) forests in G, and the circuit are the circuits. The main result of this section will be an expression for the number of proper k-colourings of G in terms of the Möbius function of L_G.

A *proper k-colouring* of G is a mapping $f : V \to \{1,\ldots,k\}$ such that $f(u) \neq f(v)$ whenver uv is an edge. If f is a mapping from V to $\{1,\ldots,k\}$, define the set $K(f)$ by

$$K(f) := \{uv \in E : f(u) = f(v)\}.$$

Note that the components of $K(f)$ form a contraction and that f is a proper colouring of G if and only if $K(f) = \emptyset$. Let $F_=(A,k)$ denote the number of mappings f from V to $\{1,\ldots,k\}$ such that $K(f) = A$ and let $F_\leq(A,k)$ be the number of mappings from V

to $\{1, \ldots, k\}$ such that $K(f) \geq A$. Since $B \subseteq K(f)$ if and only if f is constant on the components of B, we have

$$F_{\leq}(B, k) = k^{n-r(B)}.$$

As we also have

$$F_{\leq}(B, k) = \sum_{A \supseteq B} F_{=}(A, k),$$

by Möbius inversion we find that

$$F_{=}(A, k) = \sum_{B : B \supseteq A} \mu_{L_G}(A, B) k^{n-r(B)}.$$

Thus we have proved the following.

17.1 Theorem. *The number of proper k-colourings of the graph G is equal to*

$$\sum_i \left(\sum_{|A|=i} \mu_{L_G}(0, A) \right) k^{n-i}. \qquad \square$$

In other words

$$\sum_i \left(\sum_{|A|=i} \mu_{L_G}(0, A) \right) x^{n-i}$$

is the chromatic polynomial of G. We have shown that the coefficients of the chromatic polynomial are the level numbers of the broken circuit complex of G. (Perhaps we should say of L_G.)

Let Ω be a set of points in the projective space $PG(d, q)$. Then, as we noted earlier, the intersections of S with the hyperplanes of $PG(d, q)$ are the elements of a geometric lattice. The rank of a subset of Ω is equal to the dimension of the space spanned by it. We are interested in counting the number of hyperplanes of $PG(n, q)$ that contain no point of Ω. (Well, I am interested. You may have to fake it.)

The points of Ω can be represented by vectors $x_1, \ldots, x_n$ in $V(d+1, q)$. If $a \in V(d+1, q)$ then the vectors x_i such that $a^T x_i = 0$ are a hyperplane in the geometric lattice L determined by Ω. Denote the hyperplane corresponding to the vector a by $h(a)$. If $S \subseteq \Omega$, define $f(S)$ to be the number of vectors a such that $h(a) = S$ and let $g(S)$ be the number of vectors a such that $h(a) \supseteq S$. Then $g(S) = q^{d-r(S)}$ and consequently

$$f(S) = \sum_T \mu_L(S, T) q^{d-r(T)}.$$

Therefore

$$f(\emptyset) = \sum_T \mu_L(0,T)q^{d-r(T)} = \sum_{k\geq 0}\left(\sum_{r(T)=k}\mu_L(0,T)\right)q^{d-k}.$$

This is a polynomial in q, which we will denote by $F_L(q)$. It is called the *characteristic polynomial* of L.

Exercise: Show that the number of t-tuples of vectors $a_1,\dots,a_t$ such that $\cap_i h(a_i) = \emptyset$ is equal to $F_L(q^t)$.

There is coding theory view of all this, which is both interesting and useful. Suppose that we arrange the vectors $x_1,\dots,x_n$ given above into a $(d+1)\times n$ matrix, G say. The row space of G is a linear code over the field $GF(q)$. If $a \in V$ then a^TG is a code word and the weight of this word is the number of elements of Ω not in $h(a)$. Thus the hyperplanes of the lattice of flats of Ω correspond to the code words with minimal non-zero weight. Further, there is a vector a such that $h(a)$ is disjoint from Ω if and only if there is a code word with weight n, and the number of such codewords is equal to $f(\emptyset)$.

18. Points and Hyperplanes

The main result in this section is that a geometric lattice always has at least as many hyperplanes as points. The lattice of subspaces of a finite vector space shows that equality can occur. The proof makes use of another interesting result.

18.1 Theorem (Dowling and Wilson [7]). *Let L be a finite lattice. If $\mu_L(p,1)\neq 0$ for all elements p of L then there is a permutation σ of the elements of L such that $q\vee\sigma(q) = 1$ for all q in L.*

Proof. We use Lemma 6.1. Let g be the real-valued function on L defined by

$$g(p) = \begin{cases} 1 & \text{if } p = 1; \\ 0, & \text{otherwise.}\end{cases}$$

Let G be the matrix with rows and columns indexed by the elements of L and with $(G)_{pq} = g(p\vee q)$. We can complete the proof by showing that $\det G\neq 0$. By Equation (1) from dets we have

$$\det G = \prod_x \sum_y \mu_L(x,y)g(y) = \prod_x \mu_L(x,1)$$

and, by our hypothesis on L, it follows that $\det H\neq 0$. □

Any permutation σ satisfying the condition of Theorem 18.1 must map 0 to 1. Hence if L is geometric and p is a point of L then $\sigma(p)$ must be a hyperplane. Therefore σ determines an injection of the points of L into its hyperplanes, and so the number of hyperplanes in a geometric lattice is at least as large as the number of points. Actually a somewhat stronger statement can be made.

If L is a lattice let W_k denote the number of elements in L with height k. If L is geometric with height n then $W_0 = W_n = 1$, while W_1 is the number of points and W_{n-1} is the number of hyperplanes. We have just seen that Theorem 18.1 implies that $W_1 \leq W_{n-1}$. (The numbers W_i are sometimes referred to as the *Whitney numbers of the first kind.*)

18.2 Corollary. *If L is a geometric lattice with rank d then*

$$W_0 + \cdots + W_k \leq W_{d-k} + \cdots + W_d.$$

Proof. Assume L is geometric and $p \vee \sigma(p) = 1$ for all elements p of L. Since

$$r(p \vee \sigma(p)) + r(p \wedge \sigma(p)) \leq r(p) + r(\sigma(p))$$

we see that if $r(p) + r(\sigma(p)) \geq d$ for all p. So if $r(p) \leq k$ then $r(\sigma(p)) \geq d - k$. $\square$

A well known conjecture asserts that for a geometric lattice the numbers W_i form a unimodal sequence. The previous corollary is probably the best evidence for this conjecture. As stated, Corollary 18.2 is due to Dowling and Wilson [7], but the most interesting case is when $k = 1$, where the result was first established by Basterfield and Kelly [3], and independently by C. Greene [8]. If equality holds in Corollary 18.2 then Dowling and Wilson [7] prove that L must be modular; in the case $k = 1$ this was also observed in [3, 8]. We take this further in p̂t-hyps.

If p is a point in L then $\sigma(P)$ must be a complement to p. (This can be viewed as a consequence of the fact that points are modular elements.) Thus it is natural to ask if there could be a permutation σ such that $\sigma(p)$ is a complement of p, for all elements p in the lattice. This can be achieved under suitable conditions.

18.3 Theorem (Dowling [6]). *Let L be a lattice such that $\mu_L(0,p)\mu_L(p,1) \neq 0$, for any element p. Then there is a permutation σ of L such that $\sigma(p)$ is a complement of p, for all p in L.*

Proof. Let $G(p)$ denote the set of all elements x of L' such that $x \vee p < 1$. Let M be the matrix with rows and columns indexed by L, such that

$$(M)_{pq} := \mu(G(p)_{\leq q}).$$

By Lemma 10.1, the pq-entry of M is zero if p and q are not complements, so we can prove the theorem by showing that $\det M \neq 0$. (Perhaps it is worth noting that M is probably not symmetric.)

If $x \in L'$ then $p \neq 0$ and so $\sum_{z \leq p} \mu_L(0, z) = 0$. Hence

$$\begin{aligned} 0 &= \sum_{z \leq q,\ z \vee p < 1} \mu_L(0, z) + \sum_{z \leq q,\ z \vee p = 1} \mu_L(0, z) \\ &= \mu(G(p)_{\leq q}) + \sum_{z \leq q,\ z \vee p = 1} \mu_L(0, z). \end{aligned}$$

Therefore

$$\mu(G(p)_{\leq q}) = - \sum_{z \leq q, z \vee p = 1} \mu_L(0, z).$$

If H denotes the matrix we used in the proof of Theorem 18.1 and D is the diagonal matrix with $(D)_{pp} = \mu_L(0, p)$ then

$$M = -Z^T D H,$$

from which the theorem follows immediately. □

Unfortunately Theorem 18.3 does not seem to lead to any strengthening of Corollary 18.2.

19. Modular Lattices

We have seen that if L is geometric with height n then $W_1 \leq W_{n-1}$. It is reasonable to ask what can be said if equality holds. As we will see in the next section, the answer is that L must be modular. For this to make sense we must first define modular lattices themselves.

We start with an identity due to Dedekind, which holds in any lattice L. Suppose that a, b and c are elements of L. Then $a \vee (b \wedge c)$ lies below both $a \vee b$ and $a \vee c$. Hence $a \vee (b \wedge c) \leq (a \vee b) \wedge (a \vee c)$ and so we see we have proved that, if $a \leq c$ then

$$a \vee (b \wedge c) \leq (a \vee b) \wedge c. \tag{19.1}$$

This is Dedekind's identity. A lattice is *modular* if equality holds in Dedekind's identity for all a, b and c in L with $a \leq b$. It is not hard to verify that any sublattice of a modular lattice is modular, and that products of modular lattices are modular. The dual of a modular lattice is modular. All modular lattices are ranked, but the proof of this is left to you as well.

Exercise: Show that a lattice is modular if and only if it is both upper and lower semi-modular.

A word of warning here. We defined modular elements of geometric lattices. Note though that even if a is a modular element of L, neither $[0,a]$ nor $[a,1]$ need be modular.

Our next lemma shows that we can test if a lattice is modular without looking at all triples a, b and c where $a \leq c$.

19.1 Lemma. *A lattice L is modular if and only if $a \vee (b \wedge c) = c$ whenever $c \in [a, a \vee b]$.*

Proof. Maybe this will be left for you too. □

We now turn to characterisations of modular geometric lattices.

19.2 Lemma. *Let L be a geometric lattice. An element b of L is modular if and only if $a \vee (b \wedge c) = (a \vee b) \wedge c$ whenever $a \leq c$.*

Proof. Suppose that $a \leq c$ and let w and w' respectively denote $(a \vee b) \wedge c$ and $a \vee (b \wedge c)$. Note that $w' \geq w$, by Dedekind's identity. We have

$$b \wedge c \geq b \wedge (a \vee (b \wedge c)) \geq b \wedge (b \wedge c) = b \wedge c$$

and

$$b \vee a \leq b \vee ((a \vee b) \wedge c) \leq b \vee (b \vee a),$$

therefore $b \wedge w' = b \wedge c$ and $b \vee w = b \vee a$. It is even easier to verify that $b \wedge w = b \wedge c$ and $b \vee w' = b \vee a$. So if b is modular we have both

$$r(w) + r(b) = r(b \vee w) + r(b \wedge w) = r(b \vee a) + r(b \wedge c)$$

and

$$r(w') + r(b) = r(b \vee w') + r(b \wedge w') = r(b \vee a) + r(b \wedge c),$$

whence $r(w) = r(w')$. As $w \leq w'$, this implies that $w = w'$.

Now suppose that $a \vee (b \wedge c) = (a \vee b) \wedge c$ whenever $a \leq c$. We show that no two complements of b are comparable. But if a and c are complements to b and $a \leq b$ then

$$a = a \vee 0 = a \vee (b \wedge c) = (a \vee b) \wedge c = 1 \wedge c = c.$$

Hence the complements of b form an antichain, and so b is a modular element. □

19.3 Corollary. *A geometric lattice is modular if and only if each element is modular.* □

19.4 Lemma. *A geometric lattice is modular if and only if each hyperplane is a modular element of L.*

Proof. Suppose $c \in L$ and $c \not\leq h$. Then $h \vee c = 1$ and $h \wedge c$ is a hyperplane in the interval $[0, c]$. We first show that if h is modular in L then $h \wedge c$ is modular in $[0, c]$. If h is a hyperplane and $a \not\leq h$ then

$$r(h) + r(a) \geq r(h \vee a) + r(h \wedge a) = r(1) + r(h \wedge a).$$

Hence h is modular if and only if $r(a) - r(h \wedge a) = r(1) - r(h) = 1$, i.e., if and only if $a \wedge h$ is covered by a, for any a in L such that $a \not\leq h$. If $b \leq a$ then $(h \wedge a) \wedge b = h \wedge b$. If h is modular then b must cover $h \wedge b$, whence $(h \wedge a) \wedge b$ is covered by b. It follows that $h \wedge a$ is a modular hyperplane in $[0, a]$.

Now suppose that L is a geometric lattice in which every hyperplane is modular. We prove L is modular by induction on its height. If $a < 1$ then all hyperplanes in $[a, 1]$ are modular. If $a > 0$ then, by the previous paragraph, all hyperplanes in $[0, a]$ are modular. (You should show that if a covers b then there is a hyperplane h such that $h \wedge a = b$.) By induction it follows that all proper intervals of L are modular.

Now let a and b be any two elements of L. We want to verify that

$$r(a) + r(b) = r(a \vee b) + r(a \wedge b). \tag{19.2}$$

If either $a \vee b < 1$ or $a \wedge b > 0$ then this already holds, by our induction hypothesis. We therefore assume that a and b are complements. Since every element of L is the intersection of the hyperplanes containing it, there is a hyperplane h containing a but not b. If $h \wedge b = 0$ then, since h is modular b must be a point. As all points are modular, equality then holds in (19.2).

Hence we may assume that $h \wedge b > 0$. Since $[0, h]$ is modular

$$r(a) + r(b \wedge h) = r(a \vee (b \wedge h)) + r(a \wedge b \wedge h) = r(a \vee (b \wedge h))$$

and, since $[b \wedge h, 1]$ is modular

$$r(b) + r(a \vee (b \wedge h)) = r(1) + r(b \wedge (a \vee (b \wedge h))).$$

Combining these two inequalities we deduce that

$$r(b) + r(a) = r(1) + r(b \wedge (a \vee (b \wedge h))) - r(b \wedge h).$$

Let w denote $b \wedge (a \vee (b \wedge h))$. Clearly $w \geq b \wedge h$, if we can show that $w = b \wedge h$ then $r(a) + r(b) = 1$, as required.

Since $a \wedge b = 0$, it is trivial to verify that $b \wedge h$ is a complement to a in $[0, a \vee (b \wedge h)]$. Further

$$a \wedge w = a \wedge (b \wedge (a \vee (b \wedge h))) \leq a \wedge b = 0$$

while

$$a \vee (b \wedge h) = a \vee (a \vee (b \wedge h)) \geq a \vee (b \wedge (a \vee (b \wedge h))) \geq a \vee (b \wedge h)$$

and so w is a second complement to a in $[0, a \vee (b \wedge h)]$. Since this interval is contained in $[0, h]$ it is modular and, as $b \wedge h \leq w$ it follows that $w = b \wedge h$. □

A *line* in a geometric lattice is an element of height two, i.e., the join of two points.

19.5 Lemma. *A hyperplane h in a geometric lattice is modular if and only if $h \wedge \ell > 0$, for every line ℓ.*

Proof. Let h be a hyperplane. To show h is modular we need only verify that if $b \not\leq h$ then b covers $b \wedge h$. Assume by way of contradiction that $b \not\leq h$ and b does not cover h. We will use this to find a line meeting h in 0.

Let c be a complement to $b \wedge h$ in $[0, b]$. Since $r(b \wedge c) = 0$ we have $r(c) \geq r(b) - r(b \wedge h)$, consequently $r(c) \geq 2$ and so there is a line ℓ lying below c. But then

$$h \wedge \ell = h \wedge (b \wedge \ell) = (h \wedge b) \wedge \ell = 0.$$

It follows that h is modular. □

20. Points and Hyperplanes (again)

We will prove now that if the number of points in a geometric lattice L is equal to the number of hyperplanes then L is modular. Both proofs proceed by showing that if L has rank d and $W_1(L) = W_{d-1}(L)$ then any hyperplane meets any line non-trivially. (So there was some point to the trials of the previous section.) One reason this result is so interesting is that the structure of complemented modular lattices is very restricted: every complemented modular lattice is a direct sum of subspace lattices $\mathcal{B}_q(n)$ and copies of $\mathcal{B}(1)$.

20.1 Theorem (Basterfield and Kelly [3]). *Let L be a geometric lattice of rank d. Then $W_1(L) = W_{d-1}(L)$ if and only if L is modular.*

Proof. The proof that $W_1(L) = W_{d-1}(L)$ when L is modular is left as an exercise; we will only consider the sufficiency of the stated condition.

Let G be the 01-matrix with rows and columns indexed by L and with $(G)_{ab} = 1$ if and only if $a \vee b = 1$. As we saw in dets, we have

$$G = ZFZ^T$$

where $Z = Z_L$ and F is the diagonal matrix with with $(F)_{aa} = \mu(a, 1)$. Since L is geometric, F is invertible and therefore G is invertible. We claim that if $a \wedge b = 0$ then $(G^{-1})_{ab} \neq 0$. In fact we have $G^{-1} = (Z^T)^{-1}F^{-1}Z^{-1}$ and therefore

$$(G^{-1})_{ab} = \sum_{x \leq a \wedge b} \frac{\mu(x,a)\mu(x,b)}{\mu(x,1)}.$$

When $a \wedge b = 0$ this implies that

$$(G^{-1})_{ab} = \frac{\mu(0,a)\mu(0,b)}{\mu(0,1)} \neq 0.$$

We may write G in partitioned form as

$$G = \begin{pmatrix} 0 & M \\ N & X \end{pmatrix}$$

where the rows of M are indexed by the points of L and its columns by the hyperplanes. Since G is invertible the rows of M are linearly independent; this proves again that $W_1(L) \leq W_{d-1}(L)$.

Now assume that $W_1(L) = W_{d-1}(L)$. Then M and N are square invertible matrices and accordingly

$$G^{-1} = \begin{pmatrix} -N^{-1}XM^{-1} & N^{-1} \\ M^{-1} & 0 \end{pmatrix}.$$

What matters here is the zero submatrix of G^{-1}—its presence shows that if a is not a point and h is a hyperplane of L then $(G^{-1})_{ah} = 0$. This implies that $h \wedge a > 0$ and consequently for all hyperplanes h and all lines ℓ of L we have $h \wedge \ell > 0$. Therefore L is modular by Lemma 19.5. □

The proof of Theorem 20.1 can be extended to show that if equality holds in Corollary 18.2 then L is modular. (This is not an unreasonable exercise.) If L is a modular geometric lattice then it can be shown that the Whitney numbers $W_i(L)$ form a symmetric unimodal sequence. As we noted earlier, it has been conjectured that this sequence is unimodal for any geometric lattice. There is one more thing that can be proved.

20.2 Corollary (Greene [8]). *Let L be a geometric lattice of rank d. Then $W_1(L) \leq W_i(L)$, and if equality holds then $i = d - 1$.*

Proof. Let L be a geometric lattice and let f be the map from L to L defined by

$$f(x) = \begin{cases} x, & \text{if } r(x) < k; \\ 1, & \text{otherwise.} \end{cases}$$

Then f is order-preserving and its image is a geometric lattice (albeit, not a sublattice of L). The image of L under f is an *upper truncation* of L. Suppose that $W_1(L) = W_i(L)$ and let L' be the geometric lattice obtained by truncating L at height $i + 1$. Then L' is geometric and we have

$$W_1(L) = W_1(L') \leq W_i(L') = W_i(L).$$

Further, if the first and last terms here are equal then L' is modular, by the previous theorem.

Assume $i < d$ and choose an element a of L with rank $i - 2$. Then the interval $[a, 1]$ in L has height four and therefore it contains a set of four independent points. These four points generate a sublattice isomorphic to $\mathcal{B}(4)$, and the elements of rank two in it have rank i in L. Hence the interval $[a, 1]_{L'}$ has height three and contains four points and six lines. This implies that it is not modular, but every interval in a modular lattice is modular and therefore L' cannot be modular. □

Our proof of Theorem 20.1 was based on the approach of Dowling and Wilson. We now present a version of the original proof of Basterfield and Kelly. (It is simple and elegant—our only criticism is that it does not use the Möbius function.)

Assume that L is a geometric lattice with rank d. We aim to prove by induction on d that $W_1(L) \leq W_{d-1}(L)$, with equality implying that L is modular. Let p be a point and h a hyperplane of L such that $p \not\leq h$. We make two claims:

(a) $W_{d-2}[p, 1] \geq W_1[0, h]$ and
(b) $W_1[p, 1] \geq W_1[0, h]$.

To prove (a), suppose that a and b are covered by h. If $p \vee a = p \vee b$ then

$$p \vee a = p \vee a \vee b = p \vee h,$$

but since p is modular $r(p \vee a) < r(p \vee h)$ and therefore the map $x \to x \vee p$ is in injection from the hyperplanes of $[0, h]$ into the hyperplanes of L on p. Thus $W_{d-2}[0, h] \leq W_{d-2}[0, p]$.

Since $[0,h]$ is geometric, $W_1[0,h] \leq W_{d-2}[0,h]$ by induction and thus (a) is proved. For (b), the map $x \to x \vee p$ is a bijection from the points of $[0,h]$ to the lines of L on p which intersect h nontrivially.

Now we prove that $W_1(L) \leq W_{d-1}(L)$ and that, if equality holds, $W_{d-2}[p,1] = W_1[0,h]$ for any point p and hyperplane h such that $p \wedge h = 0$. Let $\bar{p}$ and $\bar{h}$ denote $W_{d-2}[p,1]$ and $W_1[0,h]$ respectively. By (a) above, $\bar{p} \geq \bar{h}$. Let n be the number of points and m be the number of hyperplanes in L. Then

$$n = \sum_p \frac{m-\bar{p}}{m-\bar{p}} = \sum_{p,h:p\wedge h=0} \frac{1}{m-\bar{p}} \geq \sum_{p,h:p\wedge h=0} \frac{1}{m-\bar{h}} = \sum_h \frac{n-\bar{h}}{m-\bar{h}}. \tag{20.1}$$

If $m < n$ then

$$\frac{n-\bar{h}}{m-\bar{h}} > \frac{n}{m}$$

whence the last term in (20.1) is strictly great than n. Hence we conclude that $m \geq n$ and, if equality holds, $\bar{p} = \bar{h}$ for any point p and hyperplane h with $p \wedge h = 0$.

Assume that $m = n$ and that p is a point and h is a hyperplane not on p. Since $[p,1]$ is geometric, we may use (b) above to deduce that

$$\bar{h} = W_1[0,h] \leq W_1[p,1] \leq W_{d-2}[1,p] = \bar{p}.$$

This implies that $W_1[0,h] = W_1[p,1]$, and our proof of (b) then implies that every line on p meets h non-trivially. Thus we have shown that if h is a hyperplane and ℓ is a line in L then $h \wedge \ell > 0$ and therefore L is modular, by Lemma 19.5.

21. Kung Fu?

We will describe some important work of J. Kung [10], in a formulation communicated privately to the author by C. Greene. This provides yet another approach to some of the work in $\widehat{\text{h}}$ypers and $\widehat{\text{pt}}$-hyps.

If f is a function on a lattice L, let $\hat{f}$ be defined by

$$\hat{f}(a) = \sum_{x \leq a} f(x).$$

Our main theorem can be viewed as providing one answer to the following problem. Suppose that A and B are subsets of a lattice L. What conditions on A and B guarantee that any function f on L with support in A is determined by the restriction $\hat{f} \restriction B$ of $\hat{f}$ to B? (Admittedly this appears to be a convoluted problem, with little hope of a useful answer arising.)

21.1 Theorem (Kung [10]). *Let A and B be subsets of the lattice L such that, if $x \in L$ then either*
(a) $x \in B$, or
(b) there exists x^* in L such that $\mu(x, x^*) \neq 0$ and $a \vee x \neq x^*$ if $a \in A$.
Then $\hat{f} \restriction B$ determines $\hat{f}$, and there is an injection $\varphi : A \mapsto B$ such that $\varphi(a) \geq a$ for all a in A.

Before embarking on the proof of this result, we present one application. Let L be a geometric lattice with rank d, let A be the set of elements of L with rank at most k and let B be the set of elements with rank at least $d-k$. If $x \in L$, let $x^* = 1$.

If $x \notin B$ then $\mu_L(x, 1) \neq 0$. If, further, $a \in A$ then

$$r(x \vee a) \leq r(x) + r(a) \leq d - k - 1 + k < d = r(1).$$

Hence the conditions of Theorem 20.1 are (and can be!) satisfied. What may we conclude? Let f_a be the function on L defined by

$$f_a(x) = \begin{cases} 1, & \text{if } x = a; \\ 0, & \text{otherwise.} \end{cases}$$

Then $\hat{f}_a(b) = 1$ if $b \geq a$, and is zero otherwise. The theorem implies that, for functions f supported on A, the linear mapping

$$f \to \hat{f} \restriction B$$

is injective. This implies that $\dim \mathbf{R}^B \geq \dim \mathbf{R}^A$, from which it follows that $|A| \leq |B|$. This provides another proof of Corollary 18.2. In fact a stronger statement can be made. The function $\hat{f}_a$ can be identifed with the row of Z_L indexed by a and therefore Kung's theorem implies that the submatrix of Z_L with rows indexed by elements of A and columns by elements of B has linearly independent rows. Hence there is an injection $\varphi : A \to B$ such that $\varphi(a) \geq a$, for all a in A.

We start the proof of Theorem 20.1 now. If x and y are elements of L such that $x < y$ and f is a function on L, we have

$$\sum_{t \in [x,y]} \mu_L(t, y) \hat{f}(t) = \sum_{t \in [x,y]} \sum_{s \leq t} \mu_L(t, y) f(s) = \sum_{s} f(s) \sum_{t \in [x \vee s, y]} \mu_L(t, y)$$
$$= \sum_{s : s \vee x = y} f(s).$$

Now suppose $x \notin B$ and let y be x^*. Then, if the support of f is contained in A, the last term above is zero and so

$$\mu_L(x, x^*)\hat{f}(x) = -\sum_{x<t\leq x^*} \mu_L(t, x^*)\hat{f}(t). \tag{21.1}$$

Condition (b) of the theorem implies that 1 must lie in B. By (21.1), if $x \notin B$ then $\hat{f}(x)$ is determined by the value of $\hat{f}$ on elements t in L such that $t > x$. It follows by induction that $\hat{f}$ is determined by $\hat{f}\restriction B$. This completes the proof.

We describe a second application of Theorem 21.1, to modular lattices. An element a in a lattice L is *join-irreducible* if, whenever $x \vee y = a$, either $x = a$ or $y = a$. In other words, a is not the join of two smaller elements. The set of all join-irreducible elements of L will be denoted by $J(L)$. Similarly we may define *meet-irreducible* elements; the set of meet-irreducible elements of L will be denoted by $M(L)$. Note $0 \in J(L)$ and $1 \in M(L)$, hence these subsets are not empty. In a geometric lattice, $J(L)$ consists of 0 and all the points, while $M(L)$ consists of the hyperplanes and 1.

Assume L is modular with $A = J(L)$ and $B = M(L)$. If $x \notin B$, define x^* to be the join of the elements which cover x. Then $[x, x^*]$ is a modular point-lattice, therefore it is geometric and $\mu_L(x, x^*) \neq 0$. Suppose a is join-irreducible. Since L is modular, the intervals $[a \wedge x, a]$ and $[x, x \vee a]$ are isomorphic (by Lemma 13.1). But this implies that $x \vee a$ is join-irreducible in $[x, x \vee a]$, which is a geometric lattice. Therefore $x \vee a$ must x, and hence cannot be equal to x^*. Thus the conditions of Kung's theorem are satisfied, whence we conclude that in a modular lattice $|J(L)| \leq |M(L)|$. As L^{op} is modular if L is and

$$J(L^{\mathrm{op}}) = M(L), \quad M(L^{\mathrm{op}}) = J(L)$$

it follows that $|J(L)| = |M(L)|$ for modular lattices. (This is a famous result of Dilworth's.)

22. Contraction and Deletion

Let L be a point lattice with point set Ω and suppose $p \in \Omega$. We define a function f from L into itself as follows:

$$f(a) = \vee\{q : q \in \Omega \setminus p,\ q \leq a, \},$$

with the understanding that $f(0) = 0$. It is easy to check that f is order preserving and that $f(x)$ is either x itself or the unique element covered by x and not in $[p, 1]$. Hence f is a decreasing map. Note that f is an order preserving and decreasing map from $L' \setminus p$ into itself.

Let M be the poset of fixed points of f. This is a join semi-lattice with a 0- and 1-element. (The latter is usually the 1-element of L.) If $M' := M \setminus \{0, 1_L\}$ then, by Lemma 8.4, M' is a retract of $L' \setminus p$ and therefore $\mu(M') = \mu(L' \setminus p)$. On the other hand by Lemma 14.1 we have

$$\mu(L') = \mu(L' \setminus p) + \mu(L'_{<p})\mu(L'_{>p}).$$

Since p is a point, $\mu(L'_{<p}) = -1$ and therefore

$$\mu_L(0,1) = \mu(M') - \mu_L(p,1).$$

There are two cases to be considered. If the join of the points of L distinct from p is equal to 1_L then $\mu(M') = \mu_M(0,1)$. If the join of the points of L distinct from p is not equal to 1_L then, because of the careful way we defined it, M' has a 1-element and $\mu(M') = 0$. We will call a point p a *co-loop* if the join h of the points distinct from p is not equal to 1. (If L is geometric and $h \neq 1$ then h is a modular hyperplane.) Since M is a semi-lattice it gives rise naturally to a lattice that we will denote by $L \backslash p$. (This is not a particularly good choice of notation, but will do for now.) We can summarise our conclusions as follows.

22.1 Lemma. *Let L be a point lattice and let p be a point of L. Then*

$$\mu_L(0,1) = \begin{cases} -\mu_L(p,1), & \text{if } p \text{ is a co-loop;} \\ \mu_{L\backslash p}(0,1) - \mu_L(p,1), & \text{otherwise.} \end{cases}$$

If L is the lattice of contractions of a graph G and e is an edge in G then e is a point in L and $L \setminus e$ is the lattice of contractions of the graph $G \setminus e$, obtained by deleting e from G. The interval $[e, 1]$ in L is the lattice of contractions of G/e, which is the graph obtained from G be contracting the edge e. It can be shown that if L is geometric then so is $L \setminus p$.

Exercise: Use Lemma 22.1 to prove the broken circuit theorem (Theorem 15.1).

The only significant application of Lemma 22.1 I know of is to geometric lattices. There are many other classes of point lattices though—the face lattices of convex polytopes, for example.

23. Null Designs

Let P be a poset. A function of *strength* at least t on P is an function f with values in some ring such that such that

$$\sum_{x \geq a} f(x) = 0$$

for any element a of P with height at most t. (In practice we assume that the ring is the ring of integers.) The most important case is when P is the lattice of all subsets of a set V, when a function of strength at least t is sometimes called a null t-design. Our basic problem is to find good lower bounds on the support of a function of strength t.

If f is any function on P and $\hat{f}$ is as in the previous section then f has strength at least t if and only if the support of $\hat{f}$ contains no elements of height t or less. Let P be a poset with zero. If $b \in P$ let f_b be the function obtained by Möbius inversion on $[0, b]$ to $\hat{f}$. Then f_b is a function on $[0, b]$, which we extend to a function on P by setting $f_b(x) = 0$ if $x \not\leq b$. It is immediate that f_b is a function of strength at least t with support contained in $[0, b]$. We have two formulas for computing the values of f_b.

23.1 Lemma. *Let P be a meet semi-lattice. If $b \in P$ then*

$$f_b(c) = \sum_{x \wedge b = c} f(x).$$

Proof. We have

$$\begin{aligned} f_b(c) = \sum_{y \leq b} \mu(c, y) \hat{f}(y) &= \sum_{y \leq b} \mu(c, y) \sum_{x \geq y} f(x) \\ &= \sum_{x, y : c \leq y \leq x} \mu(c, y) f(x) \\ &= \sum_{x} \left(\sum_{y : c \leq y \leq x} \mu(c, y) \right) f(x) \\ &= \sum_{x \wedge b = c} f(x). \end{aligned}$$

The next result is trivial to verify.

23.2 Lemma. *Let f be a function on the poset P with strength at least t and let b be an element of P in the support of $\hat{f}$ with minimal height. If $c \leq b$ then*

$$f_b(c) = \mu(c, b) \hat{f}(b).$$

□

As an immediate corollary of Lemma 23.2, we see that the support of f is bounded below by

$$|\{c \leq b : \mu(c,b) \neq 0\}|.$$

This bound can be improved in two cases. The previous two lemmas combine to yield that

$$\mu(c,b)\hat{f}(b) = \sum_{x \wedge b = c} f(x). \tag{23.1}$$

23.3 Lemma. *Let P be a meet semi-lattice and let f be a $(0, \pm 1)$-valued function of strength t on P. If b is an element of height t in P such that $\hat{f}(b) \neq 0$ then the support of f has size at least*

$$\sum_{c \leq b} |\mu(c,b)|.$$

Proof. If f is $(0, \pm 1)$-valued then $|\hat{f}(b)| \geq 1$ and (23.1) implies that there at least $|\mu(c,b)|$ elements x in P such that $x \wedge b = c$ and $f(x) \neq 0$. □

Both the previous lemma and the following theorem seem to have appeared first in unpublished work of Cho.

23.4 Theorem. *Let P be a semi-lattice, let f be a function on P with strength t which is supported by elements of height $t+1$. If b is an element of P with height $t+1$ such that $\hat{f} \neq 0$ then the support of f has size at least*

$$\sum_{c \leq b} |\mu(c,b)|.$$

If equality holds then f is $(0, \pm 1)$-valued.

Proof. Assume that, amongst all elements of height $t+1$ in the support of $\hat{f}$, we have chosen b so that $|\hat{f}(b)|$ is maximal. Our constraint on the support of f implies that if b has height $t+1$ then $\hat{f}(b) = f(b)$. By (23.1),

$$\mu(c,b) = \sum_{x \wedge b = c} \frac{f(x)}{f(b)} \leq \sum_{\substack{x : x \wedge b = c \\ f(x) \neq 0}} 1.$$

The theorem follows at once. □

It is reasonable to ask which meet semi-lattices we would like to apply the results of this section to. The first two cases of interest are $\mathcal{B}(n)$ and $\mathcal{B}_q(n)$, which are lattices. We also have the subspace lattice of a polar space. Finally suppose that V is a d-dimensional vector space over a finite field and U is a subspace of V. Then the set of subspaces of V which intersect U in the zero subspace is a meet semi-lattice. Both these last examples have the property that any interval is the subspace lattice of a projective space, and thus its Möbius function is known.

For $\mathcal{B}(n)$, it is not too hard to prove that a function of strength at least t has support of size at least 2^{t+1}. This is stronger that we have just proved. But for $\mathcal{B}_q(n)$ the results of this section are the strongest known. It is sometimes possible to give a simpler expression for

$$\sum_{c \le b} |\mu(c,b)|.$$

If $P = \mathcal{B}(n)$ and b is a set of size $(t+1)$, this sum equals 2^{t+1}. If P is $\mathcal{B}_q(n)$ and b has dimension $t+1$ then it equals

$$\prod_{i=0}^{t}(1+q^i).$$

If $P = \mathcal{P}(n)$ and b is a partition with exactly $n-k$ cells then our sum equals $(n-k)!$. (These claims all follow from [1: Proposition 4.20].)

Appendix: Posets and Lattices

The material in Aigner [1: Chapter II] provides more than enough background for our purposes. Crawley and Dilworth's book [5] is a masterpiece.

Let $(P, \leq)$ be a poset. We usually use the same symbol to denote both the poset and its underlying set, i.e, we will usually be lazy and write P rather than $(P, \leq)$. If $a, b \in P$ and $a \leq b$ but $a \neq b$, we write $a < b$. If $a < b$ and there is no element of P strictly between a and b, we say that *a covers b*. We define $(P^{\mathrm{op}}, \preceq)$ to be the poset with on the same set of elements as P, with $a \preceq b$ if and only if $b \leq a$. We say P^{op} is obtained from P be *reversal.* The interval $[a, b]$ in P consists of all elements x such that $a \leq x \leq b$.

Suppose x, y and a belong to P. We call a the *least upper bound* of x and y in P if, whenever $z \in P$ and $x, y \leq z$, we have $a \leq z$. We denote it by $x \vee y$. Similarly we define the *greatest lower bound*, and denote it by $x \wedge y$. For a general poset neither $x \vee y$ nor $x \wedge y$ need exist; if they are defined for all pairs x and y then P is a lattice. Note that every lattice has an underlying partial order, and that this partial order determines the lattice completely. If $x \wedge y$ is defined for all pairs x and y then we may define $x \vee y$ to be

$$\vee\{a \in P : a \geq x, y\}.$$

This definition may not work if P is infinite, but such concerns are above us.

A *point* in a lattice is an element which covers 0. (The term atom may also be met.) A complement of a in a lattice is an element x such that

$$a \vee x = 0, \quad a \wedge x = 1.$$

A lattice is *complemented* if every element has a complement, and is *relatively complemented* when all intervals in it are complemented. The Boolean lattice $\mathcal{B}(n)$ is relatively complemented, as is the lattice of subspaces of a finite-dimensional vector space.

References

[1] M. Aigner, *Combinatorial Theory* (Springer, Berlin) 1979.

[2] K. Baklawski, Galois connections and the Leray spectral sequence, *Advances in Math.* **25** (1977), 191–215.

[3] J. G. Basterfield and L. M. Kelly, A characterisation of n points which determine n hyperplanes, *Proc. Camb. Phil. Soc.* **64** (1968), 585–588.

[4] A. Björner and J. W. Walker, A homotopy formula for partially ordered sets, *Europ. J. Combinatorics* **4** (1983), 11–19.

[5] P. Crawley and R. P. Dilworth, *Algebraic Theory of Lattices.* (Prentice Hall, Englewood Cliffs) 1973.

[6] T. A. Dowling, Complementing permutations in finite lattices, *J. Combinatorial Theory, Series B* **23** (1977), 223-226.

[7] T. A. Dowling and R. M. Wilson, Whitney number inequalities for geometric lattices, *Proc. Amer. Math. Soc.* **47** (1975), 504–512.

[8] C. Greene, A rank inequality for geometric lattices, *J. Combinatorial Theory* **9** (1970), 357–364.

[9] C. Greene, The Möbius function of a partially ordered set, in *Ordered Sets.* Edited by I. Rival (Reidel, Dordrecht) 1981, pp. 555–581.

[10] J. P. S. Kung, Matchings and Radon transforms in lattices I. Consistent lattices, *Order*, **2** (1985), 105–112.

[11] B. Lindström, Determinants on semilattices. *Proc. Amer. Math. Soc.*, **20** (1969), 207–208.

[12] L. Lovász, *Combinatorial Problems and Exercises.* North-Holland, Amsterdam, 1979.

[13] D. Quillen, Homotopy properties of the poset of nontrivial p-subgroups of a group, *Advances in Math.* **28** (1978), 101–128.

[14] G.-C. Rota, On the foundations of combinatorial theory. I: Theory of Möbius functions, *Z. Wahrsch. Verw. Gebiete* **2** (1964), 340–368.

[15] R. P. Stanley, *Enumerative Combinatorics, Vol. I.* (Wadsworth, Monterey) 1986.

[16] J. W. Walker, Homotopy type and Euler characteristic of partially ordered sets, *Europ. J. Combinatorics* **2** (1981), 373–384.

[17] H. Whitney, A logical expansion in mathematics, *Bull. Amer. Math. Soc.* **38** (1932), 572–579.

[18] H. S. Wilf, Hadamard determinants, Möbius functions and the chromatic number of a graph, *Bull. Amer. Math. Soc.*, **74** (1968), 960–964.