Möbius functions of lattices

Andreas Blass
Department of Mathematics
University of Michigan
Ann Arbor, MI 48109-1003
ablass@umich.edu

and

Bruce E. Sagan
Department of Mathematics
Michigan State University
East Lansing, MI 48824-1027
sagan@math.msu.edu

December 1, 1995

Key Words: bounded below set, broken circuit, dominance order, lattice, Möbius function, non-crossing partition, semimodular, shuffle poset, supersolvable, Tamari lattice

AMS subject classification (1991): Primary 06A07; Secondary 06B99, 05A17, 05A18.

Proposed running head:

Möbius functions of lattices

Send proofs to:

Bruce E. Sagan
Department of Mathematics
Michigan State University
East Lansing, MI 48824-1027

Tel.: 517-355-8329
FAX: 517-336-1562
Email: sagan@mth.msu.edu

Abstract

We introduce the concept of a bounded below set in a lattice. This can be used to give a generalization of Rota's broken circuit theorem to any finite lattice. We then show how this result can be used to compute and combinatorially explain the Möbius function in various examples including non-crossing set partitions, shuffle posets, and integer partitions in dominance order. Next we present a generalization of Stanley's theorem that the characteristic polynomial of a semimodular supersolvable lattice factors over the integers. We also give some applications of this second main theorem, including the Tamari lattices.

1 Bounded below sets

In a fundamental paper [25], Whitney showed how broken circuits could be used to compute the coefficients of the chromatic polynomial of a graph. In another seminal paper [20], Rota refined and extended Whitney's theorem to give a characterization of the Möbius function of a geometric lattice. Then one of us [21] generalized Rota's result to a larger class of lattices. In this paper we will present a theorem for an arbitrary finite lattice that includes all the others as special cases. To do so, we shall need to replace the notion of a broken circuit by a new one which we call a bounded below set. Next we present some applications to lattices whose Möbius functions had previously been computed but in a less simple or less combinatorial way: shuffle posets [13], non-crossing set partition lattices [15, 19], and integer partitions under dominance order [5, 6, 12].

The second half of the paper is dedicated to generalizing the result of Stanley [23] that the characteristic polynomial of a semimodular supersolvable lattice factors over the integers. We replace both supersolvability and semimodularity by weaker conditions which we call left-modularity and the level condition, respectively, in such a way that the conclusion still holds. Examples of this factorization (not covered by Stanley's theorem) are provided by certain shuffle posets and the Tamari lattices [9, 14]. We end with a further generalization of Rota's theorem and a section of comments and questions.

Throughout this paper L will denote a finite lattice. Any relevant definitions not given can be found in Stanley's text [24]. We will use $\wedge$ for the meet (greatest lower bound) and $\vee$ for the join (least upper bound) in L. Since L is finite it also has a unique minimal element $\hat{0}$ and a unique maximal element $\hat{1}$. The *Möbius function of* L, $\mu : L \to \mathbf{Z}$, is defined recursively by

$$\mu(x) = \begin{cases} 1 & \text{if } x = \hat{0}, \\ -\sum_{y<x} \mu(y) & \text{if } x > \hat{0}. \end{cases}$$

We let $\mu(L) = \mu(\hat{1})$. Note that μ is the unique $\mathbf{Z}$-valued function on L such that $\sum_{y \leq x} \mu(y) = \delta_{\hat{0}x}$ (Kronecker delta).

Our goal is to give a new combinatorial description of $\mu(x)$. Let $A(L)$ be the set of atoms of L, i.e., those elements covering $\hat{0}$. Give $A(L)$ an arbitrary partial order, which we denote $\trianglelefteq$ to distinguish it from the partial order $\leq$ in L. So $\trianglelefteq$ can be anything from a total order to the total incomparability order induced by $\leq$. A nonempty set $D \subseteq A(L)$ is *bounded below* or *BB* if, for every $d \in D$ there is an $a \in A(L)$ such that

$$\begin{aligned} a &\vartriangleleft d \text{ and} &(1)\\ a &< \bigvee D. &(2) \end{aligned}$$

So a is simultaneously a strict lower bound for d in the order $\trianglelefteq$ and for $\bigvee D$ in $\leq$. We will say that $B \subseteq A(L)$ is *NBB* if B does not contain any D which is bounded below. In this case we will call B an *NBB base* for $x = \bigvee B$. We can now state our main result.

Theorem 1.1 *Let L be any finite lattice and let $\trianglelefteq$ be any partial order on $A(L)$. Then for all $x \in L$ we have*

$$\mu(x) = \sum_B (-1)^{|B|} \tag{3}$$

where the sum is over all NBB bases B of x and $|\cdot|$ denotes cardinality.

Proof. For $x \in L$ define $\tilde{\mu}(x) = \sum_B (-1)^{|B|}$ summed over all NBB bases B of x. To prove that $\tilde{\mu}(x) = \mu(x)$ it suffices to show $\sum_{y \leq x} \tilde{\mu}(y) = \delta_{\hat{0}x}$. If $x = \hat{0}$ then $x = \bigvee B$ only for $B = \emptyset$ which is NBB. So

$$\sum_{y \leq \hat{0}} \tilde{\mu}(y) = \tilde{\mu}(\hat{0}) = (-1)^0 = 1$$

as desired.

If $x > \hat{0}$ then to get $\sum_{y \leq x} \tilde{\mu}(y) = 0$ we first set up a corresponding signed set $\mathcal{S}$. Let

$$\mathcal{S} = \{B \ : \ B \text{ an NBB base for some } y \leq x\}.$$

The sign of $B \in \mathcal{S}$ will be $\epsilon(B) = (-1)^{|B|}$. Then from the definitions

$$\sum_{y \leq x} \tilde{\mu}(y) = \sum_{B \in \mathcal{S}} \epsilon(B).$$

If we can find a sign-reversing involution on $\mathcal{S}$ then this last sum will be zero and we will be done. From the set of atoms $a \leq x$ pick one, a_0, which is minimal with respect to $\trianglelefteq$. Consider the map on $\mathcal{S}$ defined by $\iota(B) = B \triangle a_0$ where $\triangle$ is symmetric difference. (Here and below we omit set braces around singletons writing, for example, a_0 instead of $\{a_0\}$.) Clearly ι is a sign-reversing involution provided it is well defined, i.e., we must check that B being NBB implies that $\iota(B)$ is NBB.

If $\iota(B) = B \setminus a_0$ then clearly $\iota(B)$ is still NBB. We will do the case $B' := \iota(B) = B \cup a_0$ by contradiction. Suppose $B' \supseteq D$ where D is bounded below. Then $a_0 \in D$ since B itself is NBB. Let a be the corresponding element guaranteed by the definition of a bounded below set. Then $a \vartriangleleft a_0$ and $a < \bigvee B' \leq x$ which contradicts the definition of a_0. ∎

Here is an example to illustrate this result. Suppose the lattice L and partial order $\trianglelefteq$ are as given in Figure 1. To find the bounded below sets, note that

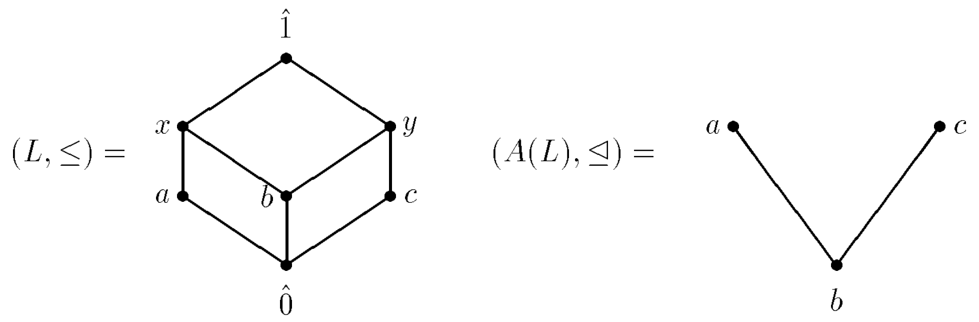

Figure 1: A lattice L and partial order $\trianglelefteq$ on $A(L)$

by (1) no set containing an element minimal in $\trianglelefteq$ is BB. Furthermore (1) and (2) together imply that no single element set is BB either. (These observations will be important in our other examples.) Thus the only possible BB set is $\{a, c\}$, and it does satisfy the definition since $b \triangleleft a, c$ and $b < \bigvee\{a, c\} = \hat{1}$. So x has one NBB base, namely $\{a, b\}$ and so $\mu(x) = (-1)^2$ which is easily checked from the definition of μ. Similarly $\mu(y) = (-1)^2$. Finally $\hat{1}$ has no NBB bases and so $\mu(\hat{1}) = 0$ (the empty sum).

We should see why our theorem implies Rota's broken circuit result. To do this we need to recall some definitions. Let L be a geometric lattice with rank function ρ. It is well known, and easy to prove, that if $B \subseteq A(L)$ then $\rho(\bigvee B) \leq |B|$. Call B *independent* if $\rho(\bigvee B) = |B|$ and *dependent* otherwise. A *circuit* C is a minimal (with respect to inclusion) dependent set. Now let $\trianglelefteq$ be any total order on $A(L)$. Then each circuit C gives rise to a *broken circuit* $C' = C \setminus c$ where c is the first element of C under $\trianglelefteq$. A set $B \subseteq A(L)$ is *NBC* (no broken circuit) if B does not contain any broken circuit and in this case B is an *NBC base* for $x = \bigvee B$. Rota's NBC theorem [20] is as follows.

Theorem 1.2 (Rota) *Let L be a finite geometric lattice and let $\trianglelefteq$ be any total order on $A(L)$. Then for all $x \in L$ we have*

$$\mu = (-1)^{\rho(x)} \cdot (\textit{number of NBC bases of } x).$$

To derive this result from Theorem 1.1, we first prove that when L is geometric and $\trianglelefteq$ is total then the NBB and NBC sets coincide. For this it suffices to show that every broken circuit is bounded below and that every bounded below set contains a broken circuit. If $C' = C \setminus c$ is a broken circuit, then for every $c' \in C'$ we have $c \triangleleft c'$ and $c < \bigvee C = \bigvee C'$ so C' is BB. For the other direction, if D is bounded below then consider the $\trianglelefteq$-first element d of D and let $a \in A(L)$ be the element guaranteed by the BB definition. Then by (2) we have $\rho(\bigvee D \vee a) < |D \cup a|$. So

$D \cup a$ is dependent and contains a circuit C. Now (1) and the choice of a and d show that for the corresponding broken circuit we have $C' \subseteq D$. Thus NBB and NBC sets are the same in this setting. Finally Rota's expression for $\mu(x)$ is obtained from ours by noting that when L is semimodular then all NBB bases for x have the same size, namely $\rho(x)$. Similar arguments show that the main result of [21] is a special case of Theorem 1.1.

Another corollary of this theorem is a special case of Rota's Crosscut Theorem [20].

Theorem 1.3 (Rota) *Let L be any finite lattice. Then for all $x \in L$*

$$\mu(x) = \sum_{B} (-1)^{|B|} \tag{4}$$

where the sum is over all $B \subseteq A(L)$ such that $\bigvee B = x$.

To obtain (4) it suffices to take $\trianglelefteq$ to be the total incomparability order in Theorem 1.1. In fact our theorem (where the partial order on $A(L)$ is arbitrary) can be viewed as interpolating between the Crosscut Theorem (where $A(L)$ forms an antichain) and the NBC Theorem (where $A(L)$ forms a chain).

This raises the question of why one would want to consider an arbitrary partial order $\trianglelefteq$ on $A(L)$ when one can always take the one induced by the order in L. The reason is that the number of terms in the sum (4) is generally much larger than the number in (3). From the viewpoint of efficient computation of μ, the best scenario is the same as the one in the geometric case where (3) has exactly $|\mu(x)|$ terms, all of the same sign. A partial order $\trianglelefteq$ on $A(L)$ for which this happens for all $x \in L$ will be called *perfect*. There are posets where no such $\trianglelefteq$ exists, such as the k-equal intersection lattices which were introduced by Björner, Lovász and Yao [3, 4]. However, all the examples we will consider in the next sections are perfect. Another thing to note is that if $\trianglelefteq$ is perfect then so is any linear extension of it. However, to make the combinatorics of μ as clear as possible it is often best to take a perfect $\trianglelefteq$ with the least possible number of order relations.

2 Non-crossing partitions

The non-crossing partition lattice was first studied by Kreweras [15] who showed its Möbius function to be a Catalan number. By using NBB sets we can combinatorially explain this fact and relate these bases to the standard NBC bases for the ordinary partition lattice.

If π is a partition of $[n] := \{1, 2, \ldots, n\}$ into k subsets, or blocks, then we write $\pi = B_1/\ldots/B_k \vdash [n]$. When it will cause no confusion, we will not explicitly write

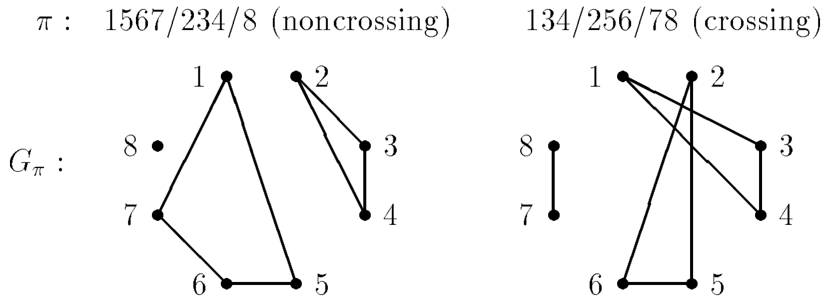

Figure 2: Partitions and their graphs

out any blocks that are singletons. The set of $\pi \vdash [n]$ form a lattice Π_n under the refinement ordering. We say that π is *non-crossing* if there do not exist $i, k \in B$ and $j, l \in C$ for two distinct blocks B, C of π with $i < j < k < l$. Otherwise π is *crossing*. The set of non-crossing partitions of $[n]$ forms a meet-sublattice NC_n of Π_n with the same rank function. However unlike Π_n, NC_n is not semimodular in general since if $\pi = 13$ and $\sigma = 24$ then $\pi \wedge \sigma = \hat{0}$ and $\pi \vee \sigma = 1234$ so

$$\rho(\pi) + \rho(\sigma) = 2 < 3 = \rho(\pi \wedge \sigma) + \rho(\pi \vee \sigma).$$

Here and throughout this paper, semimodularity refers to upper-semimodularity.

Another way to view non-crossing partitions will be useful. Let $G = (V, E)$ be a graph with vertex set $V = [n]$. Then G is *non-crossing* if there do not exist edges $ik, jl \in E$ with $i < j < k < l$. Equivalently, G is non-crossing if, when the vertices are arranged in their natural order clockwise around a circle and the edges are drawn as straight line segments, no two edges of G cross geometrically. Given a partition π we can form a graph G_π by representing each block $B = \{i_1 < i_2 < \ldots < i_l\}$ by a cycle with edges $i_1i_2, i_2i_3, \ldots, i_li_1$. (If $|B| = 1$ or 2 then B is represented by an isolated vertex or edge, repectitively.) Then it is easy to see that π is non-crossing as a partition if and only if G_π is non-crossing as a graph. In Figure 2 we have displayed some partitions and their graphs.

The atoms of NC_n are the partitions of the form $\delta = ij$ where we will always assume $i < j$. Then G_δ is a single edge, so we can consider any $B \subseteq A(NC_n)$ as a graph G_B with an edge for each $\delta \in B$. Define $ij \lhd i'j'$ if and only if $j < j'$. So the poset $(A(NC_n), \unlhd)$ will be ranked with the elements at rank $j - 2$ being all atoms of the form ij. An example for $n = 4$ is given in Figure 3

In order to characterize the NBB sets, we first need a lemma.

Lemma 2.1 *Let δ, δ' be a pair of atoms such that either δ and δ' have the same rank in $(A(NC_n), \unlhd)$ or the graph of $D = \{\delta, \delta'\}$ is crossing. Then D is BB.*

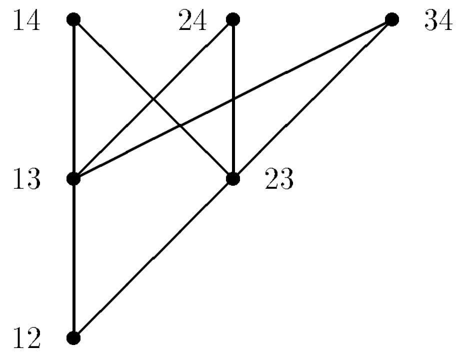

Figure 3: The partial order for $A(NC_4)$

Proof. Suppose first that the two atoms have the same rank. Then $\delta = ij$ and $\delta = i'j$ where without loss of generality $i < i' < j$. Let $\alpha = ii'$. Then $\delta, \delta' \trianglerighteq \alpha$ and $\delta \vee \delta' = ii'j > \alpha$ so D is BB.

Suppose instead that the given atoms are crossed. So we have $\delta = ij$ and $\delta = i'j'$ with $i < i' < j < j'$. Letting $\alpha = ii'$ we get the same inequalities as before, noting that $\delta \vee \delta' = ii'jj'$. ∎

Theorem 2.2 *The NBB bases of $\hat{1}$ in NC_n are all B obtained by picking exactly one element from each rank of $(A(NC_n), \trianglelefteq)$ so that the corresponding graph G_B is non-crossing.*

Proof. First suppose that B is an NBB base of $\hat{1}$. Then by the previous lemma we know that B contains at most one element from each rank and that G_B is non-crossing. If we do not pick an element from some rank then G_B is not connected. But such a non-crossing graph has a block of $\bigvee B$ for each component of G_B, contradicting $\bigvee B = \hat{1}$.

Conversely, suppose B is picked according to the two given rules. Then G_B is connected and so $\bigvee B = \hat{1}$. If $B \supseteq D$ with D a BB set we will derive a contradiction. Let $\delta \in D$ be $\trianglelefteq$-minimal and let $\alpha = ij$ be the corresponding element guaranteed by the definition of bounded below. Then (1) and our choice of δ shows that for any $\delta' = i'j' \in D$ we have $i, j < j'$. Since B is non-crossing, so is D and thus $\bigvee D$ is the same in NC_n and Π_n. It follows from (2) that there is a path in G_D

of the form $i = i_0, i_1, \ldots, i_l = j$ where each edge $\{i_k i_{k+1}\}$ is an atom of D. But our remarks about δ' imply that $i_0 < i_1$ and $i_{l-1} > i_l$ so there must be an index m such that $i_{m-1} < i_m > i_{m+1}$. Thus D has two elements from the same $\trianglelefteq$-rank, the promised contradiction. ■

Note that the graphs G_B in the previous theorem are certain spanning trees on the vertex set $[n]$. Furthermore, to get the NBB bases for all elements of NC_n it suffices to use the non-crossing restriction but picking at most one element from each rank. Finally, if one removes the non-crossing restriction one gets exactly the standard NBC bases for the geometric lattice Π_n.

It is now easy to compute the Möbius function of NC_n. It suffices to do this for $\hat{1}$ since for any $\pi = B_1/\ldots/B_k \in NC_n$, the interval $[\hat{0}, \pi] \cong \prod_i NC_{|B_i|}$. Recall that the *Catalan numbers* are defined by

$$C_n = \frac{1}{n+1}\binom{2n}{n}.$$

Corollary 2.3 (Kreweras) *We have*

$$\mu(NC_n) = (-1)^{n-1} C_{n-1}.$$

Proof. All trees on n vertices have $n-1$ edges. Furthermore it is easy to see that the number T_n of non-crossing trees on $[n]$ of the given form and C_{n-1} satisfy the same initial conditions and recurrence relation,

$$T_n = \begin{cases} 1 & \text{if } n = 1 \\ \sum_{0<i<n} T_i T_{n-i} & \text{if } n > 1. \end{cases}$$

The result now follows from Theorems 1.1 and 2.2. ■

There is another natural ordering of $A(NC_n)$ for which a result similar to Theorem 2.2 holds, namely $ij \trianglelefteq i'j'$ if and only if $[i,j] \supseteq [i',j']$ as intervals of integers. With this ordering, NBB bases of $\hat{1}$ again correspond to trees with vertex set $[n]$, but this time they are the non-crossing trees in which each vertex is either greater than all its neighbors or less than all its neighbors. It follows easily that, in any such tree, $\{1, n\}$ is an edge, and deletion of this edge leaves two smaller such trees with vertex sets $[k]$ and $[k+1, n]$ for some k. This observation easily implies that the numbers of such trees satisfy the recurrence for the Catalan numbers. In fact, there is a simple bijection between these trees and proper parenthesizations P of the product of n factors (one of the most familiar interpretations of the Catalan numbers). To describe the bijection, identify the n factors with the elements of $[n]$, in order. For any particular parenthesization P, consider the sub-products defined by P. For example, if $n = 5$ and P is $(((12)3)(45))$, then the sub-products are $(12), ((12)3), (45)$, and $(((12)3)(45))$ itself. Now to build the corresponding T, take the vertex set $[n]$ and draw, for each sub-product of P, an edge from the first to the last element of the subproduct. In the previous example, the edges would be $12, 13, 45$, and 15. It is a variant of this ordering which allows us to compute μ for the non-crossing B_n and D_n lattices.

3 Non-crossing B_n and D_n

In this section, we apply Theorem 1.1 to calculate the Möbius invariants $\mu(\hat{1})$ of the non-crossing B_n and D_n (and intermediate) lattices. These lattices were introduced by Reiner [19] who computed μ using generating functions. Non-crossing B_n consists of those partitions π of $\{1, 2, \ldots, n, -1, -2, \ldots, -n\}$ that satisfy three conditions: First, π is invariant under the involution $k \mapsto -k$. Second, at most one block of π is fixed by this involution; if there is such a block, it is called the *zero-block* of π. Third, the partition is non-crossing (as in Section 2), with respect to the ordering $1 < 2 < \ldots < n < -1 < -2 < \ldots < -n$. (The first two of these conditions determine a lattice isomorphic to that associated with the hyperplane arrangement B_n.) Non-crossing D_n is the subposet consisting of those π for which the zero-block, if present, does not consist of only a single pair $\{k, -k\}$. An intermediate lattice can be associated to every subset $S \subseteq [n]$, by allowing the zero-block to be $\{k, -k\}$ only if $k \notin S$. We use the notations NCB_n, NCD_n, and $NCBD_n(S)$ for these lattices; thus $NCB_n = NCBD_n(\emptyset)$ and $NCD_n = NCBD_n([n])$.

We begin by calculating $\mu(\hat{1})$ for NCB_n and afterward indicate the minor changes needed to handle the rest of these lattices. The atoms of NCB_n are of three sorts. First, there are the partitions where one block consists of two positive numbers, say $\{i, j\}$, another block is $\{-i, -j\}$, and all the remaining blocks are singletons. Second, there are the partitions where one block consists of a positive and a negative number, say $\{i, -j\}$, another is $\{-i, j\}$, and the rest are singletons. Third, there are the partitions whose only non-singleton block is of the form $\{i, -i\}$. Following Zaslavsky [26], we depict atoms as signed edges in a graph whose vertex set is $[n]$. An atom of the first sort is depicted as a positive edge ij, one of the second sort is depicted as a negative edge ij, and one of the third sort is depicted as a (negative) half-edge at i. (There is no such thing as a positive half-edge.) To avoid confusion, we emphasize that these signed graphs are quite different from the graphs with vertex set $\{1, 2, \ldots, n, -1, -2, \ldots, -n\}$ used in deciding whether a partition is non-crossing. We sometimes identify an atom with the corresponding edge; in particular, we may refer to an atom as positive or negative.

We partially order the atoms as follows. Associate to each atom, depicted as a signed edge ij, the interval $[i, j] \subseteq [n]$; in the case of a half-edge at i, the interval consists of just i. Then define $a \lhd b$ to mean that either the interval associated to atom a properly includes that associated to atom b or the two intervals are equal and a is negative while b is positive. (Notice that, apart from signs and half-edges, this matches the ordering described at the end of Section 2.)

Regarding atoms as signed edges, we regard sets of atoms as signed graphs with vertex set $[n]$. It is not difficult (though a bit tedious) to verify that every NBB set B has, as a graph, the following properties.

(i) No two signed edges cross.

(ii) No vertex i has a neighbor $j < i$ and also a neighbor $j' > i$.

(iii) Any path joining two vertices i and j and consisting entirely of vertices k with $i \leq k \leq j$ must be just a single signed edge ij.

(iv) There is no cycle of length ≥ 3. (A 2-cycle, consisting of a positive edge and a negative edge in the same place, is permitted; we refer to such a pair of signed edges as a double edge.)

(v) No two negative atoms in B have disjoint associated intervals.

(vi) If the interval associated to one atom in B is properly included in the interval associated to another, and if the former atom is negative, then so is the latter.

(vii) If there is a negative atom in B, then there is exactly one, say a, that is $\lhd$-maximal (i.e., its interval is inclusion-minimal); all other atoms in B are negative if their intervals properly include that of a and positive otherwise.

(viii) B has at most one half-edge, has at most one double-edge, and cannot have both.

Actually, only items (i), (ii), (v), and (vi) in this list directly use the NBB assumption. The other four items follow from these purely graph-theoretically.

Conversely, any signed graph satisfying (i) through (viii) is NBB when viewed as a set of atoms of NCB_n. We leave the verification to the reader, with the hint that items (i), (v), and (vi) ensure that the join of any atoms from this set is the same whether computed in NCB_n or in B_n, because none of these atoms cross when regarded as partitions of $\{1, 2, \ldots, n, -1, -2, \ldots, -n\}$.

From this characterization of the NBB sets in NCB_n, we easily obtain a characterization of the NBB bases of $\hat{1}$. These bases (regarded as signed graphs) are obtainable as follows. First, take a non-crossing tree T with vertex set $[n]$ in which each vertex is either greater than all its neighbors or less than all its neighbors. (This part is just as at the end of Section 2.) Then pick either an edge or a vertex of T. If you picked an edge e, then make it a double edge, give negative signs to all the edges of T whose intervals properly include that of e, and give all remaining edges of T positive signs. If you picked a vertex v, then attach a (negative) half-edge at v, give negative signs to all edges of T whose interval contains v, and give all remaining edges of T positive signs.

Finally, to apply Theorem 1.1, we count the NBB bases for $\hat{1}$, i.e., we count the signed trees of the sort just described. We already saw in Section 2 that there are

$$C_{n-1} = \frac{1}{n}\binom{2n-2}{n-1} = \frac{1}{2n-1}\binom{2n-1}{n}$$

ways to choose T. Then there are $2n-1$ ways to choose an edge or vertex, since there are n vertices and $n-1$ edges. After this choice, the rest of the construction of the NBB signed graph is completely determined. So the number of NBB bases for $\hat{1}$ is

$$C_{n-1} \cdot (2n-1) = \binom{2n-1}{n}.$$

Every NBB base of $\hat{1}$ has exactly n elements, namely the $n-1$ edges of T (with signs) plus either an extra edge if you chose an edge and doubled it or an extra half-edge if you chose a vertex. Therefore, $\mu(NCB_n) = (-1)^n\binom{2n-1}{n}$.

The calculation for $NCBD_n(S)$ is almost exactly the same. The only difference is that half-edges can occur only at vertices not in S. Thus, in the description of NBB bases for $\hat{1}$, we need only replace "pick an edge or a vertex of T" with "pick an edge of T or a vertex not in S." Thus, the number of options at this step is no longer $2n-1$ but only $2n-1-|S|$. Therefore, we obtain, in agreement with Reiner's calculation ([19]), the Möbius invariant for $NCBD_n(S)$:

$$\mu(NCBD_n(S)) = (-1)^n C_{n-1} \cdot (2n-1-|S|).$$

4 Shuffle Posets

The poset of shuffles was introduced by Greene [13]. We need to recall some of his definitions and results before applying Theorem 1.1. Let $\mathcal{A}$ be a set, called the *alphabet of letters*. A *word over* $\mathcal{A}$ is a sequence $\mathbf{u} = u_1u_2\ldots u_n$ of elements of $\mathcal{A}$. All of our words will consist of distinct letters and we will sometimes also use $\mathbf{u}$ to stand for the set of letters in the word, depending upon the context. A *subword of* $\mathbf{u}$ is $\mathbf{v} = u_{i_1}\ldots u_{i_l}$ where $i_1 < \ldots < i_l$. If $\mathbf{u}, \mathbf{v}$ are any two words then the *restriction of* $\mathbf{u}$ *to* $\mathbf{v}$ is the subword $\mathbf{u_v}$ of $\mathbf{u}$ whose letters are exactly those of $\mathbf{u} \cap \mathbf{v}$. A *shuffle of* $\mathbf{u}$ *and* $\mathbf{v}$ is any word $\mathbf{s}$ such that $\mathbf{s} = \mathbf{u} \uplus \mathbf{v}$ as sets (disjoint union) and $\mathbf{s_u} = \mathbf{u}$, $\mathbf{s_v} = \mathbf{v}$.

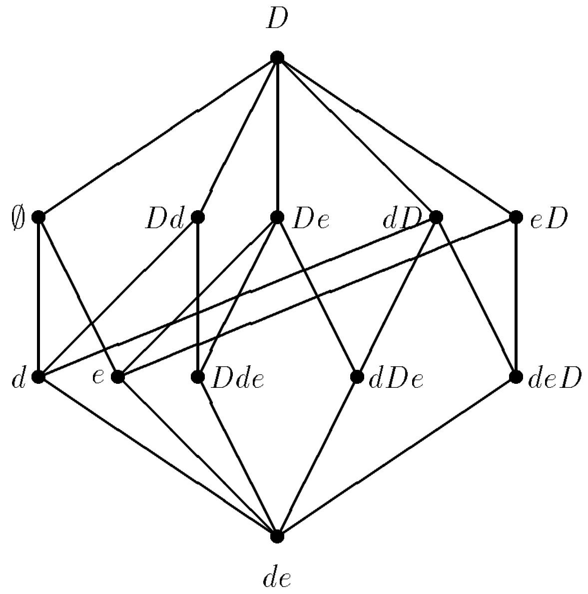

Figure 4: The lattice $\mathcal{W}_{2,1}$

Given nonnegative integers m, n Greene defined the *poset of shuffles* $\mathcal{W}_{m,n}$ as follows. Fix disjoint words $\mathbf{x} = x_1 \ldots x_m$ and $\mathbf{y} = y_1 \ldots y_n$. The elements of $\mathcal{W}_{m,n}$ are all shuffles $\mathbf{w}$ of a subword of $\mathbf{x}$ with a subword of $\mathbf{y}$. The partial order is $\mathbf{v} \leq \mathbf{w}$ if $\mathbf{v}_\mathbf{x} \supseteq \mathbf{w}_\mathbf{x}$ and $\mathbf{v}_\mathbf{y} \subseteq \mathbf{w}_\mathbf{y}$. It is easy to see that $\mathcal{W}$ has minimal element $\hat{0} = \mathbf{x}$, maximal element $\hat{1} = \mathbf{y}$ and is ranked with rank function

$$\rho(\mathbf{w}) = (m - |\mathbf{w}_\mathbf{x}|) + |\mathbf{w}_\mathbf{y}|. \tag{5}$$

For example, $\mathcal{W}_{2,1}$ is shown in Figure 4 where $\mathbf{x} = de$ and $\mathbf{y} = D$.

In order to apply the NBB Theorem we will need to describe the join operation in $\mathcal{W}_{m,n}$. Greene does this using crossed elements, not to be confused with the crossing partitions discussed in Section 2. Given $\mathbf{u}, \mathbf{v} \in \mathcal{W}_{m,n}$ then $x \in \mathbf{x}$ is *crossed in* $\mathbf{u}, \mathbf{v}$ if there exist $y_i, y_j \in \mathbf{y}$ such that $i \leq j$ and x appears before y_i in one of the two words but after y_j in the other. For example if $\mathbf{x} = def$ and $\mathbf{y} = DEF$ then for the pair $\mathbf{u} = dDEe, \mathbf{v} = Fdef$ the only crossed letter is d. The join of $\mathbf{u}, \mathbf{v}$ is then the unique word $\mathbf{w}$ greater than both $\mathbf{u}, \mathbf{v}$ such that

$$\begin{aligned} \mathbf{w}_\mathbf{x} &= \{x \in \mathbf{u}_\mathbf{x} \cap \mathbf{v}_\mathbf{x} \ : \ x \text{ is not crossed}\} \\ \mathbf{w}_\mathbf{y} &= \mathbf{u}_\mathbf{y} \cup \mathbf{v}_\mathbf{y}. \end{aligned} \tag{6}$$

In the previous example, $\mathbf{u} \vee \mathbf{v} = DEFe$. This example also shows that $\mathcal{W}_{m,n}$ is not geometric since the semimodularity law is violated:

$$\rho(\mathbf{u}) + \rho(\mathbf{v}) = 3 + 1 < 5 = \rho(\mathbf{u} \vee \mathbf{v}) \leq \rho(\mathbf{u} \vee \mathbf{v}) + \rho(\mathbf{u} \wedge \mathbf{v}).$$

The set atoms $A_{m,n} = A(\mathcal{W}_{m,n})$ consists of two types. An *a-atom*, respectively *b-atom*, is one obtained from $\mathbf{x}$ by deleting a letter of $\mathbf{x}$, respectively inserting a letter of $\mathbf{y}$. Let A_a denote the set of a-atoms and similarly for A_b. Define $\trianglelefteq$ on $A_{m,n}$ to be the poset whose relations are all those of the form $\mathbf{a} \triangleleft \mathbf{b}$ with $\mathbf{a} \in A_a$ and $\mathbf{b} \in A_b$

Lemma 4.1 *Suppose* $\mathbf{b}, \mathbf{b}' \in A_b$. *If* $\mathbf{b}, \mathbf{b}'$ *have crossed elements then* $D = \{\mathbf{b}, \mathbf{b}'\}$ *is a BB set of* $(A_{m,n}, \trianglelefteq)$.

Proof. Our hypothesis and (6) show that $(\mathbf{b} \vee \mathbf{b}')_{\mathbf{x}} \subset \mathbf{x}$ (proper containment of sets). So there is an a-atom $\mathbf{a}$ with $\mathbf{a} \subseteq (\mathbf{b} \vee \mathbf{b}')_{\mathbf{x}}$ which forces $\mathbf{a} < \mathbf{b} \vee \mathbf{b}'$ in $W_{m,n}$. Since by definition $\mathbf{a} \triangleleft \mathbf{b}, \mathbf{b}'$ the element $\mathbf{a}$ satisfies the definition of a BB set. ■

The next result will characterize the NBB bases and show that the converse of the previous lemma also holds.

Theorem 4.2 *Let* $\mathbf{s}$ *be a shuffle of* $\mathbf{x}, \mathbf{y}$ *and consider*

$$B_{\mathbf{s}} = A_a \cup \{\mathbf{b} \in A_b \ : \ \mathbf{b} \leq \mathbf{s}\}.$$

Then the NBB bases of $\mathbf{y} \in \mathcal{W}_{m,n}$ *under the given partial order are exactly the* $B_{\mathbf{s}}$.

Proof. Suppose first that B is an NBB base of $\mathbf{y}$. Then for each element $y \in \mathbf{y}$ we must have a corresponding b-atom $\mathbf{b}_y$ in order to get $\bigvee B = \mathbf{y}$. In fact there must be exactly one such atom for each $y \in \mathbf{y}$ and these atoms must all lie below a shuffle $\mathbf{s}$, for otherwise B would contain a BB pair as in Lemma 4.1. It follows that $\bigvee_y \mathbf{b}_y = \mathbf{s}$. So in order to get $\bigvee B = \mathbf{y}$ we must have $A_a \subseteq B$. Thus B is of the form $B_{\mathbf{s}}$ as desired.

Conversely, consider any $B_{\mathbf{s}}$. It is easy to see that $\bigvee B_{\mathbf{s}} = \mathbf{y}$. We will show that $B_{\mathbf{s}}$ is NBB by contradiction. Suppose that $D \subseteq B_{\mathbf{s}}$ is a BB set. Then (1) forces D to contain only b-atoms which in turn implies $\bigvee D \leq \mathbf{s}$. Now pick any $d \in D$ and let $a \in A_{m,n}$ be the atom guaranteed by the definition of BB. Then a is an a-atom and $a < \bigvee D \leq \mathbf{s}$. But this contradicts the fact that there are no a-atoms below any shuffle of $\mathbf{x}$ and $\mathbf{y}$. Thus $B_{\mathbf{s}}$ is an NBB base of $\mathbf{y}$ and our characterization is complete. ■

To determine the Möbius function of $\mathcal{W}_{m,n}$ it suffices to compute $\mu(\hat{1})$ since for any $\mathbf{w} \in \mathcal{W}_{m,n}$ the interval $[\hat{0}, \mathbf{w}]$ is isomorphic to a product of $\mathcal{W}_{p,q}$'s for certain $p \leq m$ and $q \leq n$.

Corollary 4.3 (Greene) *We have*

$$\mu(\mathcal{W}_{m,n}) = (-1)^{m+n}\binom{m+n}{m}.$$

Proof. The bases $B_{\mathbf{s}}$ all have size $m+n$ and the number of possible shuffles $\mathbf{s}$ is $\binom{m+n}{m}$. The result now follows from Theorems 1.1 and 4.2. ∎

5 The Dominance Order

Bogart [5] and Brylawski [6] first computed the two-variable Möbius function of the lattice of integer partitions under dominance. Subsequently Greene [12] gave two alternative ways to compute this function. Our NBB set characterization leads to a formula in Corollary 5.2 which is essentially equivalent to, but simpler than, Greene's second description of μ [12, Theorem 4.1].

We begin by reviewing the relevant definitions. A *partition* λ of n is a weakly decreasing sequence of positive integers $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_r$ whose sum is n. For any such partition λ and any non-negative integer k, we write $|\lambda|_k$ for $\sum_{i=1}^{k} \lambda_i$, where λ_i is interpreted as 0 for $i > r$, so $|\lambda|_k = n$ for all $k \geq r$. A partition λ *dominates* another partition ν (of the same n), written $\lambda \geq \nu$, if $|\lambda|_k \geq |\nu|_k$ for all k. In terms of Ferrers diagrams (in the English orientation) $\lambda \geq \nu$ means that the diagram for λ can be obtained from that of ν by moving some squares up to earlier rows.

It is well known that this ordering makes the set of partitions of n into a lattice $\mathcal{P}_n$; we review the construction of joins in $\mathcal{P}_n$ because it will be needed for our NBB calculations. A *composition* of n is like a partition except that the parts λ_i need not be in weakly decreasing order. The dominance order of compositions is defined exactly as for partitions, and the result is a lattice in which joins are easily computed since $|\lambda \vee \nu|_k = \max\{|\lambda|_k, |\nu|_k\}$. The join of two partitions, computed in this lattice of compositions, need not be a partition. However, for every composition λ there is a unique smallest partition above λ, which we call the *partition reflection* of λ. Joins of partitions in $\mathcal{P}_n$ can be computed by first forming the join in the dominance lattice of compositions and then forming the partition reflection of the result. We should point out that the partition reflection of a composition need not be simply the result of rearranging the parts into weakly decreasing order. For example, the partition reflection of $(1,3)$ is not $(3,1)$ but $(2,2)$.

The bottom element $\hat{0}$ of $\mathcal{P}_n$ is the partition $(1,1,\ldots,1)$ and the only atom is $(2,1,1,\ldots,1)$. By Theorem 1.3, these are the only elements of $\mathcal{P}_n$ where the

Möbius function has a non-zero value. But we can also consider the Möbius function of elements in any upper interval $[\beta, \hat{1}]$; indeed, the value of this Möbius function at some $\lambda \geq \beta$ is what is usually called the (two-variable) Möbius function $\mu(\beta, \lambda)$.

To calculate this Möbius function, we fix β, we describe the atoms of $[\beta, \hat{1}]$, and we calculate the joins of sets of atoms with particular attention to determining when one atom is below a join of others. Then, we describe an appropriate partial ordering $\trianglelefteq$ of the atoms, characterize its NBB sets, and use Theorem 1.1 to evaluate the Möbius function.

For the rest of this section, let $\beta = (\beta_1, \beta_2, \ldots, \beta_r)$ be a fixed but arbitrary partition of n. By a *wall* we mean a maximal sequence of at least two numbers k for which the corresponding β_k are equal. (This corresponds to a *flat* in Greene's terminology.) Thus, an interval $[i, j] \subseteq [1, r]$ is a wall if and only if $i < j$ and $\beta_{i-1} > \beta_i = \beta_j > \beta_{j+1}$. Here and below, we use the convention that $\beta_0 = \infty$ and $\beta_{r+1} = 0$, so that 1 (respectively, r) can be part of a wall if the first (respectively, last) two components of β are equal. If $[i, j]$ is a wall then we call i its top and j its bottom, the terminology being suggested by the Ferrers diagram.

The atoms of $[\beta, \hat{1}]$ are of two sorts:

(i) If $1 \leq i < r$ and neither i nor $i+1$ is in a wall (i.e., $\beta_{i-1} > \beta_i > \beta_{i+1} > \beta_{i+2}$), then β is covered by the partition α that agrees with β except that $\alpha_i = \beta_i + 1$ and $\alpha_{i+1} = \beta_{i+1} - 1$. We denote this atom α of $[\beta, \hat{1}]$ by $i + 1 \to i$, since its Ferrers diagram is obtained from that of β by moving a square from row $i + 1$ up to row i.

(ii) If $[i, j]$ is a wall, then β is covered by the partition that agrees with β except that $\alpha_i = \beta_i + 1$ and $\alpha_j = \beta_j - 1$. We denote this α by $j \to i$.

For both types of α, we refer to the set of k where $|\alpha|_k \neq |\beta|_k$ as the *critical interval* of α, I_α. (This is *not* the same as Brylawski's use of "critical.") It consists of only i for α of type (i) and it is $[i, j-1]$ for α of type (ii). Note that $|\alpha|_k = |\beta|_k + 1$ if and only if $k \in I_\alpha$ Note also that different atoms of $[\beta, \hat{1}]$ have disjoint critical intervals and

$$\biguplus_{\alpha} I_\alpha = [1, r-1] \setminus \{k : \ [i, k] \text{ or } [k+1, i] \text{ is a wall for some } i\}. \qquad (7)$$

So an element of $[1, r-1]$ lies in a critical interval unless it is the bottom of a wall or the immediate predecessor of the top of a wall.

Let us consider the join of some subset $\mathcal{B} \subseteq A([\beta, \hat{1}])$, first in the sense of compositions and then in the sense of partitions. If γ is the composition join, then from the previous paragraph

$$|\gamma|_k - |\beta|_k = \begin{cases} 0 & \text{if } k \notin X \\ 1 & \text{if } k \in X \end{cases} \tag{8}$$

where $X = \biguplus_{\alpha \in \mathcal{B}} I_\alpha$. If this γ is a partition, then it is also the join in the partition sense; otherwise, we must take its partition reflection. So we consider next how γ could fail to be a partition and how it is changed by reflection.

By (8) the components of γ differ from those of β by ± 1 or 0. It follows that any failure of γ to be a partition, i.e., any occurrence of $\gamma_k < \gamma_{k+1}$, must arise in one of two ways: $\beta_k = \beta_{k+1}$ or $\beta_{k+1} + 1$. The first possibility can be excluded, because it requires k and $k+1$ to be part of a wall, say $[i, j]$. If the critical interval $[i, j-1] \subseteq X$ then by (7) and (8) we have $\gamma_i = \beta_i + 1$, $\gamma_j = \beta_j - 1$, and γ agrees with β on the rest of $[i, j]$; if, on the other hand, $[i, j-1] \cap X = \emptyset$, then γ agrees with β on all of $[i, j]$. In either case, we cannot have $\gamma_k < \gamma_{k+1}$. So any occurrence of $\gamma_k < \gamma_{k+1}$ has $\beta_k = \beta_{k+1} + 1$, $\gamma_k = \beta_k - 1 = \beta_{k+1}$, and $\gamma_{k+1} = \beta_{k+1} + 1 = \beta_k$. Thus by (8) X must contain $k-1$ and $k+1$ but not k.

Consider the set Y obtained by adjoining to X all those numbers $k \notin X$ for which $k-1, k+1 \in X$ and $\beta_k = \beta_{k+1} + 1$. Let ν be the composition of n such that

$$|\nu|_k - |\beta|_k = \begin{cases} 0 & \text{if } k \notin Y \\ 1 & \text{if } k \in Y. \end{cases} \tag{9}$$

The preceding observations show that ν is a partition, because, by enlarging X to Y, we have corrected all the failures of γ to be a partition. Furthermore, ν clearly dominates γ. We claim that ν is the partition reflection of γ, i.e., that every partition η dominating γ also dominates ν. To see this, it suffices to show that $|\eta|_k > |\gamma|_k$ for $k \in Y - X$; but if this failed for some k, then η would fail to be a partition because $\eta_k < \eta_{k+1}$ just as for γ. Thus, ν is the join of $\mathcal{B}$ in $\mathcal{P}_n$.

We must also determine which atoms of $[\beta, \hat{1}]$ other than members of $\mathcal{B}$ are below ν. Distinct atoms have disjoint critical intervals, so the atoms we are looking for are those whose critical intervals are singletons k where $k \in Y - X$. Note that there is an atom with critical interval k if and only if either k and $k+1$ constitute a wall or neither of them belongs to a wall. The former alternative is irrelevant in the present context, since $k \in Y - X$ implies $\beta_k = \beta_{k+1} + 1$. So we need only consider the second alternative, where $\beta_{k-1} > \beta_k > \beta_{k+1} > \beta_{k+2}$.

We are ready to begin the description of the NBB bases for an arbitrary $\lambda \geq \beta$. Let $\mathcal{A}$ be the set of atoms of $[\beta, \lambda]$. If $\bigvee \mathcal{A} \neq \lambda$ then $\mu(\beta, \lambda) = 0$ by Theorem 1.3, so from now on we assume $\bigvee \mathcal{A} = \lambda$. Call an atom $b \in \mathcal{A}$ of the form $k+1 \to k$ *special* if $\beta_k = \beta_{k+1} + 1$ and $\mathcal{A}$ also contains atoms a, c equal to $k+2 \to k+1, k \to k-1$, respectively. The preceding discussion shows that $b < a \vee c$; conversely, if an atom

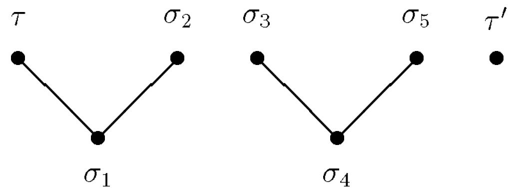

Figure 5: The partial order on an extended special run

in $\mathcal{A}$ is under the join of some other atoms in $\mathcal{A}$, then it must be special. We let $\mathcal{S} \subset \mathcal{A}$ be the set of special atoms.

Let us list the atoms in $\mathcal{A}$ according to the ordering of their critical intervals. By a *special run* in this list, we mean a maximal sequence of consecutive $\sigma \in \mathcal{S}$. Define a partial ordering $\trianglelefteq$ on $\mathcal{A}$ by imposing on each sequence

$$\tau = \sigma_0, \sigma_1, \ldots, \sigma_{q+1} = \tau' \tag{10}$$

where $\sigma_1, \ldots, \sigma_q$ is a special run (so $\tau, \tau' \notin \mathcal{S}$) the relations

$$\sigma_{3i+1} \triangleleft \sigma_{3i}, \sigma_{3i+2} \quad \text{for} \quad 1 \leq 3i+1 \leq q.$$

An example is given in Figure 5. Note that this ordering makes $\{\sigma_{3i}, \sigma_{3i+2}\}$ a BB set for each i. We let $\mathcal{S}_j$, $0 \leq j \leq 2$, be the set of elements in all special runs of the form σ_{3i+j} for some i, so $\mathcal{S} = \mathcal{S}_0 \uplus \mathcal{S}_1 \uplus \mathcal{S}_2$.

Theorem 5.1 *Let $[\beta, \lambda]$ be an interval in $\mathcal{P}_n$ with atom set $\mathcal{A}$ such that $\bigvee \mathcal{A} = \lambda$. Then λ has an NBB base in $\mathcal{A}$ if and only if there is no special run of length 1 modulo 3. If a base exists, then it is unique and equals*

$$B = \mathcal{A} \setminus \mathcal{S}_2. \tag{11}$$

Proof. If B is an NBB base of λ then it must include $\mathcal{A} \setminus \mathcal{S}$ since these elements are not $\leq$ the join of any others. Now consider any sequence of the form (10). Since $\tau \notin \mathcal{S}$ we have $\tau \in B$ and this forces $\sigma_2 \notin B$ since $\{\tau, \sigma_2\}$ is BB. Now we must put $\sigma_1, \sigma_3 \in B$ since neither is below $\bigvee(\mathcal{A} \setminus \sigma_2)$. Now repeat the argument with σ_3 in place of τ to inductively see that the only possible candidate for an NBB base is (11). But if there is a special run (10) of length $3k+1$ for some k then σ_{3k} and $\tau' = \sigma_{3k+2}$ form a BB set in B, a contradiction. If no such run exists then we never have such a pair in B and so it is NBB as desired. ∎

Corollary 5.2 *Let $[\beta, \lambda]$ be an interval in $\mathcal{P}_n$ with atom set $\mathcal{A}$. If $\bigvee \mathcal{A} \neq \lambda$ then $\mu(\beta, \lambda) = 0$. If $\bigvee \mathcal{A} = \lambda$ then let r_i, $0 \leq i \leq 2$ be the number of special runs of length i modulo 3. Then*

$$\mu(\beta, \lambda) = \begin{cases} 0 & \text{if } r_1 \geq 1 \\ (-1)^{|\mathcal{A} \setminus \mathcal{S}| + r_2} & \text{if } r_1 = 0. \end{cases}$$

Proof. We have already noted that $\mu(\beta, \lambda) = 0$ if $\bigvee \mathcal{A} \neq \lambda$. For the other two cases, suppose first that $r_1 \geq 1$. Then λ has no NBB base in $[\beta, \lambda]$ by Theorem 5.1 so $\mu(\beta, \lambda) = 0$ by Theorem 1.1. If $r_1 = 0$ then by the same two results

$$\mu(\beta, \lambda) = (-1)^{|B|} = (-1)^{|\mathcal{A} \setminus \mathcal{S}| + |\mathcal{S}_0| + |\mathcal{S}_1|}.$$

If a special run has length congruent to 0 (respectively, 2) modulo 3 then its contribution to $|\mathcal{S}_0| + |\mathcal{S}_1|$ is 0 (respectively, 1) modulo 2. The last case now follows from the previous displayed equation. ■

6 Supersolvable Lattices

Stanley [23] defined a *supersolvable lattice* to be a pair (L, Δ) where L is a lattice, $\Delta : \hat{0} = x_0 < x_1 < \ldots < x_{n-1} < x_n = \hat{1}$ is a maximal chain of L, and Δ together with any other chain of L generates a distributive lattice. One often refers to L as supersolvable, Δ being tacitly understood. It is easy to see that a supersolvable lattice has a rank function ρ. He showed that, if such an L is also semimodular, then its characteristic polynomial

$$\chi(L, t) = \sum_{x \in L} \mu(\hat{0}, x) t^{n - \rho(x)} \tag{12}$$

factors as $(t - a_1)(t - a_2) \cdots (t - a_n)$, where a_i is the number of atoms of L that are below x_i but not below x_{i-1}. Our purpose in this section is to use Theorem 1.1 to prove Stanley's factorization of the characteristic polynomial for a wider class of lattices. For this purpose we will replace both supersolvability and semimodularity by weaker hypotheses.

To state the first of our two hypotheses, we define an element x of a lattice L to be *left-modular* if, for all $y \leq z$,

$$y \vee (x \wedge z) = (y \vee x) \wedge z.$$

It is standard to call (x, z) a *modular pair* if the preceding equation is satisfied by every y that is $\leq z$. So x is left-modular if and only if every pair with x on the left is modular. Note that left-modularity, unlike the analogously defined right-modularity, is a self-dual concept of lattice theory. Our first hypothesis is that

L has a maximal chain Δ, all of whose elements are left-modular.

In this case we will call L itself *left-modular.*

Fix a maximal chain $\Delta : \hat{0} = x_0 < x_1 < \ldots < x_{n-1} < x_n = \hat{1}$ whose elements may or may not be left-modular. We partition the set A of atoms of L into pieces $A_i = \{a \in A \mid a \leq x_i \text{ but } a \not\leq x_{i-1}\}$, which we call the *levels* of A, and we partially order A by setting $a \lhd b$ if and only if a is in an earlier level than b, where "earlier" means having a smaller subscript. We say that this order $\unlhd$ is *induced* by the chain Δ. Note that an atom a cannot be $\leq \bigvee S$ for any set S of atoms from strictly lower levels, for there is an x_i that is $\geq$ all elements of S but not $\geq a$. We can now state our second hypothesis.

If $\unlhd$ is induced by Δ and $a \lhd b_1 \lhd b_2 \lhd \ldots \lhd b_k$ then $a \not\leq \bigvee_{i=1}^{k} b_i$.

A lattice L having a chain Δ with this property will be said to satisfy the *level condition.* If L is a finite lattice with a maximal chain satisfying both the left-modular and level conditions then it is called an *LL* lattice.

Proposition 6.1 *1. If L is supersolvable then it is left-modular but not conversely.*

2. If L is semimodular then it satisfies the level condition (for any maximal chain) but not conversely.

Proof. Stanley has already pointed out that first statement holds [23, Proposition 2.2 and ff.]. For the second, recall that semimodularity implies that if $x \in L$ and a is an atom not below x then $x \vee a$ covers x. Now prove the implication by induction on k, the number of b atoms in the level condition. The cases $k = 0$ and $k = 1$ are obvious, so suppose the result holds for $k - 1$ but fails for k. So we have $a \lhd b_1 \lhd b_2 \lhd \ldots \lhd b_k$ and $a \leq \bigvee_{i=1}^{k} b_i$. Let $c = \bigvee_{i=1}^{k-1} b_i$ and notice that $a \not\leq c$ by induction hypothesis. Thus $c < c \vee a \leq c \vee b_k$. But b_k is an atom, so $c \vee b_k$ covers c. Therefore, $c \vee a = c \vee b_k \geq b_k$. But this is absurd, as $c \vee a$ is a join of atoms from levels strictly earlier than that of b_k.

The implication is not reversible as the shuffle posets $\mathcal{W}_{m,1}$ serve as counterexamples. It is not hard to check that the level condition holds when $\mathbf{y} = y_1$ with Δ being the chain where the x_i's are first removed one at a time and then y_1 is added. On the other hand, as long as $m \geq 1$, there are pairs of atoms not covered by their join (atoms that add y_1 in different places), so $\mathcal{W}_{m,1}$ is not a semimodular lattice. ∎

We intend to generalize to LL lattices Stanley's factorization of the characteristic polynomial. Since LL lattices need not be ranked, we need a suitable substitute for the rank function used in defining the characteristic polynomial. Let L be a left-modular lattice, with left-modular maximal chain Δ inducing a partition into levels A_i as above. Define the *generalized rank* of $x \in L$ to be

$$\rho(x) = \text{ number of } A_i\text{'s containing atoms} \leq x.$$

Later in this section, we shall relate this ρ to lengths of chains, but for now we use it in (12) (with n still denoting the length of the chain Δ) to define the characteristic polynomial of any left-modular lattice. We shall obtain a factorization of this polynomial for any LL lattice.

We first characterize the NBB sets using the following lemma.

Lemma 6.2 *If a and b are distinct atoms from the same level A_i in a left-modular lattice, then $a \vee b$ is above some atom c from an earlier level.*

Proof. Since b is below x_i but not below x_{i-1}, we have $x_{i-1} < x_{i-1} \vee b \leq x_i$. By maximality of the chain Δ, it follows that $x_{i-1} \vee b = x_i \geq a \vee b$. Applying left-modularity of x_{i-1} with $y = b$ and $z = a \vee b$, we find that

$$b \vee (x_{i-1} \wedge (a \vee b)) = (b \vee x_{i-1}) \wedge (a \vee b) = a \vee b.$$

But $b \not\geq a$ so the right side of this equation is not b. Since the left side is also not b, we have that $x_{i-1} \wedge (a \vee b)$ is not $\hat{0}$ and so is above some atom c satisfying the lemma. ■

Theorem 6.3 *In an LL lattice, the NBB sets are exactly those subsets of A that have at most one member in each level A_i.*

Proof. The level condition immediately implies that, if $B \subseteq A$ has no two members from the same level then no subset of it is BB. Conversely, if B has two members at the same level then, by the lemma just proved, those two constitute a BB set.■

The following lemma is useful not only for our primary goal, factoring the characteristic polynomial of an LL lattice, but also for relating our generalized rank functions to lengths of chains.

Lemma 6.4 *Let B be an NBB set in an LL lattice. Then every atom $a \leq \bigvee B$ is at the same level as some element of B. In particular, any NBB base for x has exactly $\rho(x)$ elements.*

Proof. It suffices to prove the first statement, since the second follows from it and Theorem 6.3. So suppose B and a were a counterexample to the first statement. Let A_j be the level containing a, and let y be the join of all the elements of B

of higher level than A_j. Since B has no element at level A_j, we have $a \leq \bigvee B \leq x_{j-1} \vee y$. Setting $z = a \vee y$, we obtain from the left-modularity of x_{j-1} that $(y \vee x_{j-1}) \wedge z = y \vee (x_{j-1} \wedge z)$.

On the left side of this last equation, both sides of the meet are $\geq a$, and therefore so is the whole left side. On the right side, since z is a join of atoms from levels A_j and higher, the level condition tells us that z is above no atom of lower level than A_j, so no atom is below $x_{j-1} \wedge z$. Therefore $x_{j-1} \wedge z = \hat{0}$ and the right side reduces to y.

Combining these results, we have $a \leq y$. But, in view of the definition of y, this contradicts the level condition. ■

We can now prove our generalization of Stanley's theorem on semimodular supersolvable lattices.

Theorem 6.5 *If L is an LL lattice then characteristic polynomial of L factors as*

$$\chi(L,t) = \prod_{i=1}^{n}(t - |A_i|)$$

Proof. Combining Theorems 1.1 and 6.3 along with the new definition of ρ we have

$$\begin{aligned}
\chi(L,t) &= \sum_{x \in L} \sum_{\substack{B \text{ NBB} \\ \bigvee B = x}} (-1)^{|B|} t^{n-\rho(x)} \\
&= \sum_{\forall i:\ |B \cap A_i| \leq 1} (-1)^{|B|} t^{n-|B|} \\
&= \prod_{i=1}^{n}(t - |A_i|). \quad \blacksquare
\end{aligned}$$

To close this section, we point out two situations where our generalized rank can be described in terms of lengths of chains.

The first of these situations is in a semimodular (hence ranked), left-modular (hence LL) lattice L. As above, let $\trianglelefteq$ be the partial order of $A(L)$ induced by a maximal chain of left-modular elements and let ρ be the generalized rank function given by the levels induced by the same chain. Also, let $\trianglelefteq^*$ be an arbitrary linear extension of $\trianglelefteq$. We write NBB and NBB* to mean NBB with respect to $\trianglelefteq$ and $\trianglelefteq^*$, respectively. Notice that every NBB* set is also an NBB set. As noted in our discussion of Rota's NBC theorem, the size of any NBB* base for x is the ordinary rank of x. By Lemma 6.4, the size of any NBB base for x is $\rho(x)$. So the generalized rank $\rho(x)$ agrees with the ordinary rank of x provided x has at least one NBB* base. This proviso can be reformulated as $\mu(x) \neq 0$, so the characteristic polynomial is the same for both notions of rank. The proviso cannot be omitted. A chain is a distributive lattice (hence also semimodular and left-modular) in which our generalized rank is 1 for all elements except $\hat{0}$ and therefore differs from the ordinary rank if the chain has more than two members.

The second situation is described in the following proposition, which connects ρ to lengths of chains even in some unranked lattices.

Proposition 6.6 *Let L be a left-modular lattice which is atomic. Then for all $x \in L$ we have*

$$\rho(x) = \text{ the length of the longest } \hat{0} \text{ to } x \text{ chain.}$$

Proof. Let $\hat{0} = x_0 < x_1 < \ldots < x_{n-1} < x_n = \hat{1}$ be the left-modular chain used for the definitions of the sets A_i and thus of ρ. There are, for any $x \in L$, exactly $\rho(x)$ values of i such that A_i contains an atom $\leq x$. The elements $x \wedge x_i$ for these values of i (along with $\hat{0}$) constitute a chain of length $\rho(x)$ from $\hat{0}$ to x. It remains to show that no chain from $\hat{0}$ to x is longer, and for this it suffices to show that, for all $a < b$ in L, $\rho(a) < \rho(b)$.

Let $a < b$ and choose, by atomicity of L, an atom $p \leq b$ such that $p \not\leq a$ and such that $p \in A_j$ for the smallest possible j. Then, by atomicity again, $x_{j-1} \wedge b \leq a$. This and the left-modularity of x_{j-1} imply

$$a = a \vee (x_{j-1} \wedge b) = (a \vee x_{j-1}) \wedge b.$$

If there were an atom $q \in A_j$ that is $\leq a$, then we would have $x_{j-1} < q \vee x_{j-1} \leq x_j$ and, by maximality of the chain of x_i's, $q \vee x_{j-1} = x_j$. But then $(a \vee x_{j-1}) \wedge b \geq x_j \wedge b \geq p$. This and $a \not\geq p$ contradict the equation displayed above.

So there is no such q. But that means that the A_i's counted by $\rho(b)$ include all those counted by $\rho(a)$ and at least one more, namely A_j. Therefore, $\rho(a) < \rho(b)$.■

Notice that the hypotheses of the preceding proposition do not imply that L is ranked. A counterexample is given by the six-element lattice obtained from the eight-element Boolean algebra by removing two co-atoms.

7 More examples

We will now give two examples where our factorization theorem can be applied but Stanley's cannot because the lattices involved are not semimodular.

The first example is the shuffle poset $\mathcal{W}_{m,1}$ with the maximal chain

$$\mathbf{x} = x_0 < x_1 < \ldots < x_{m+1} = \mathbf{y} \tag{13}$$

where $x_1, \ldots, x_m$ are obtained by deleting the letters of $\mathbf{x}$ in some order. (Note that we are using x_i to denote an element of the chain rather than a letter of $\mathbf{x}$.) Greene [13] showed that the given chain satisfies the supersolvability condition even when extended in $\mathcal{W}_{m,n}$ by adding the letters of $\mathbf{y}$ in some order. So by Proposition 6.1 $\mathcal{W}_{m,1}$ is left-modular. We also mentioned in the proof of the same proposition that this poset satisfies the level condition. (However the level condition does not hold in $\mathcal{W}_{m,n}$ for general $n \geq 2$, and this is reflected by the fact that the corresponding characteristic polynomials usually do not factor over the integers.) It is now easy to see that the number of new atoms below x_i in (13) is

$$|A_i| = \begin{cases} 1 & \text{if } i \leq m \\ m+1 & \text{if } i = m+1. \end{cases}$$

From Theorem 6.5 we immediately get the following.

Corollary 7.1 *We have*

$$\chi(\mathcal{W}_{m,1}, t) = (t-1)^m (t-m-1). \quad \blacksquare$$

Note that $\mathcal{W}_{m,n}$ is ranked in the ordinary sense and Greene computed $\chi(\mathcal{W}_{m,n}, t)$ using the usual rank function. But, by Proposition 6.6, this rank function coincides with ours.

For our second example we will use the Tamari lattices [8, 9, 10, 11, 14]. Consider all proper parenthesizations π of the word $x_1 x_2 \ldots x_{n+1}$. It is well known that the number of these is the Catalan number C_n. Partially order this set by saying that σ covers π whenever

$$\pi = \ldots((AB)C)\ldots \qquad \text{and} \qquad \sigma = \ldots(A(BC))\ldots$$

for some subwords A, B, C. The corresponding poset turns out to be a lattice called the *Tamari lattice* T_n. Figure 6(a) gives a picture of T_3.

A *left bracket vector*, $(v_1, \ldots, v_n)$, is a vector of positive integers satisfying

1. $1 \leq v_i \leq i$ for all i and
2. if $S_i = \{v_i, v_i + 1, \ldots, i\}$ then for any pair S_i, S_j either one set contains the other or $S_i \cap S_j = \emptyset$.

The number of left bracket vectors having n components is also C_n. In fact given a parenthesized word π we have an associated left bracket vector $v = (v_1, \ldots, v_n)$ defined as follows. To calculate v_i, start at x_i in π and move left, counting the number of x's and the number of left parentheses you pass (including x_i itself) until

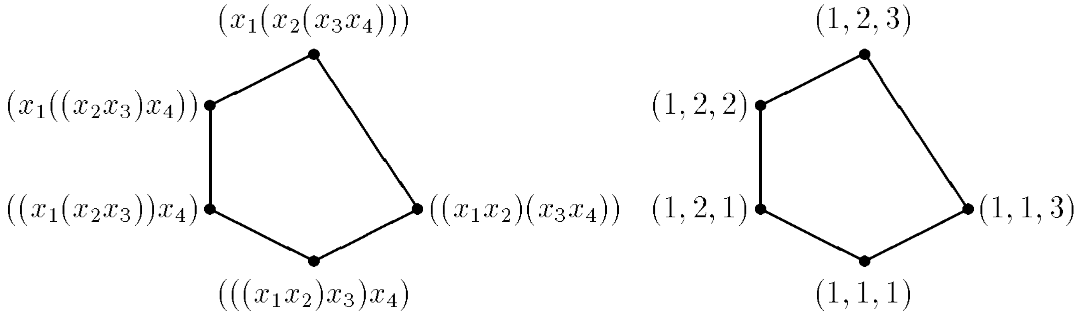

(a) Parenthesized version (b) Left bracket version

Figure 6: The Tamari lattice T_3

these two numbers are equal. Then $v_i = j$ where x_j is the last x passed before the numbers balance. It is not hard to show that this gives a bijection between parenthesizations and left bracket vectors, thus inducing a partial order on the latter. In fact this induced order is just the component-wise one. Figure 6(b) gives the bracket vector version of T_3.

Left bracket vectors are also directly related to the trees G_B described in Theorem 2.2, but with $n+1$ in place of n. Indeed, given a left bracket vector $(v_1, \ldots, v_n)$, we obtain such a tree G_B with vertex set $[n+1]$ by joining $i+1$ to v_i for $i = 1, 2, \ldots, n$, and all trees as in Theorem 2.2 can be obtained in this way.

We should note that we have used the left bracket vector rather than the more traditional right bracket one because when using the former it is easier to describe the join operation than the meet. (The situation is reversed for the right bracket vector.) This makes it simpler to work with some of our conditions which only involve joins. Expressions for the join and meet can be obtained by dualizing results in [14] and [17] respectively.

Proposition 7.2 *Given left bracket vectors $v = (v_1, \ldots, v_n)$ and $w = (w_1, \ldots, w_n)$ then*

$$v \vee w = (\max\{v_1, w_1\}, \ldots, \max\{v_n, w_n\}).$$

If we let $m = (m_1, \ldots, m_n)$ where $m_i = \min\{v_i, w_i\}$ then $v \wedge w = (l_1, \ldots, l_n)$ where the l_i are computed recursively by

$$l_i = \min\{m_i, l_{m_i}, l_{m_i+1}, \ldots, l_{i-1}\}. \quad \blacksquare$$

To show that T_m is LL, consider the chain

$$\begin{aligned}\Delta:\ & (1,\ldots,1) < (1,2,1,\ldots,1) < (1,2,2,1,\ldots,1) \\ & < (1,2,3,1,\ldots,1) < (1,2,3,2,1,\ldots,1) < \ldots < (1,2,3,\ldots,n).\end{aligned}$$

It is easy to see that this is of maximum length in T_n. The description of join in Proposition 7.2 also gives a quick proof that the level condition holds. To verify left-modularity will take more work.

A typical element of Δ looks like

$$x = (1, 2, \ldots, j-1, i, 1, \ldots, 1) \tag{14}$$

where $i \leq j$. Take $y = (y_1, \ldots, y_n) \leq z = (z_1, \ldots, z_n)$ in T_n. We wish to compute $y \vee (x \wedge z) = (c_1, \ldots, c_n)$ so we first consider $x \wedge z$. Following the notation of Proposition 7.2 we have

$$m = (z_1, \ldots, z_{j-1}, \min\{z_j, i\}, 1, \ldots, 1).$$

Using the recursive construction on the first $j-1$ components of m leaves them unchanged since these are the initial components of a vector in T_{j-1}. Also the last $n-j$ components will still be 1 because m_i is in the minimum taken for l_i. So $x \wedge z$ is the same as m except possibly in the jth component which is z_j if $z_j \leq i$ or $\min\{i, z_i, \ldots, z_{j-1}\}$ if $z_j \geq i$. Since $y \leq z$ we have $y_j \leq z_j$ for all j and so

$$y \vee (x \wedge z) = (z_1, \ldots, z_{j-1}, c_j, y_{j+1}, \ldots, y_j) \tag{15}$$

where

$$c_j = \begin{cases} z_j & \text{if } z_j \leq i \\ \max\{y_j, \min\{i, z_i, \ldots, z_{j-1}\}\} & \text{if } z_j \geq i. \end{cases} \tag{16}$$

Now we concentrate on $(y \vee x) \wedge z = (d_1, \ldots, d_n)$. Clearly

$$y \vee x = (1, 2, \ldots, j-1, \max\{i, y_j\}, y_{j+1}, \ldots, y_n)$$

Taking the minimum vector to compute $(y \vee x) \wedge z$ we get

$$m' = (z_1, \ldots, z_{j-1}, d'_j, y_{j+1}, \ldots, y_n)$$

where d'_j is z_j if $z_j \leq i$ and $\max\{i, y_j\}$ if $z_j \geq i$. So in $(y \vee x) \wedge z$ we have $d_i = z_i = c_i$ for $i < j$ by the same reasoning as before. We claim that $d_i = y_i$ for $i > j$. From the minimization procedure in Proposition 7.2 we have $d_i \leq m_i = y_i$. But in any lattice $(y \vee x) \wedge z \geq y \vee (x \wedge z)$ and so

$$d_i \geq c_i \text{ for all } 1 \leq i \leq n. \tag{17}$$

Since $c_i = y_i$ for $i > j$ this forces the desired equality.

To complete the proof of left-modularity we must show $d_j = c_j$ and this breaks up into three cases. If $z_j \leq i$ then we have that $d_j = z_j = c_j$. If $y_j \leq i \leq z_j$ then

$d_j' = i$ and $d_j = \min\{i, z_i, \ldots, z_{j-1}\}$. Comparison of this last equation with (16) gives $d_j \leq c_j$ and then (17) results in $d_j = c_j$. Finally if $i \leq y_j \leq z_j$ then $d_j' = y_j$ and

$$d_j = \min\{y_j, z_{y_j}, \ldots, z_{j-1}\} \leq y_j = c_j.$$

Using (17) again finishes this last case.

It is now an easy matter to compute the characteristic polynomial. Note that the atoms of T_n are all elements of the form $(1, \ldots, 1, j, 1, \ldots, 1)$ with the $j \geq 2$ in the jth position. If x is as in (14) then it covers no new atoms if $i < j$ and exactly one new atom if $i = j$. Using Theorem 6.5 this translates as follows.

Corollary 7.3 *We have*

$$\chi(T_n, t) = t^{\binom{n-1}{2}}(t-1)^{n-1}. \quad \blacksquare$$

8 A further generalization

In this section, we briefly describe a generalization of Theorem 1.1. As before, we work with a finite lattice L, but instead of an additional partial ordering of the atoms we use an arbitrary function M assigning to each $x \in L \setminus \hat{0}$ a non-empty set $M(x)$ of atoms $\leq x$. For comparison with Theorem 1.1, the reader should regard the ordering $\trianglelefteq$ used there as inducing the function

$$M(x) = \{a \in A(L) \mid a \text{ is } \trianglelefteq\text{-minimal among atoms } \leq x\}.$$

But what follows applies also to M's that do not arise from partial orderings in this way.

For any set B of atoms, let $S(B)$ be the subset obtained by deleting from B all members of $M(x)$ for all $x \geq \bigvee B$. By the *core* of B, we mean the set obtained by starting with B and repeatedly applying S until the decreasing sequence $B, S(B), S(S(B)), \ldots$ stabilizes; the final $S^n(B)$ is the core of B. (This can also be described as the largest subset of B unchanged by S.) We call B *coreless* if its core is empty.

Theorem 8.1 *Let L be any finite lattice and let M be any function assigning to every $x \in L \setminus \hat{0}$ a nonempty set $M(x)$ of atoms that are $\leq x$. Then for all $x \in L$ we have*

$$\mu(x) = \sum_B (-1)^{|B|}$$

where the sum is over all coreless $B \subseteq A(L)$ whose join is x.

We omit the proof since it is essentially the same as that of Theorem 1.1. The only difference is that, instead of choosing a $\trianglelefteq$-minimal atom $a_0 \le x$, we choose an $a_0 \in M(x)$.

Notice that, if M happens to be obtained from a partial order $\trianglelefteq$ as described above, then non-empty cores with respect to M are the same as BB sets with respect to $\trianglelefteq$, and therefore coreless sets with respect to M are the same as NBB sets with respect to $\trianglelefteq$.

Notice also that the theorem of this section is "stronger," in the sense of having fewer summands in the formula for the Möbius function, when $M(x)$ is smaller, for this will make $S(B)$ larger and therefore make B less likely to be coreless. From this point of view, one should always take $M(x)$ to be a singleton. Of course considerations of naturality or clarity may make other choices of M preferable (just as they may make nonlinear orderings $\trianglelefteq$ preferable to the more efficient linear orderings in applications of Theorem 1.1).

9 Comments and Questions

(1) It would be interesting to find other applications of our two main theorems. The higher Stasheff-Tamari posets as recently defined in [8] are obvious candidates. However, one would first have to prove that they are lattices.

(2) In a previous paper [21] one of us proved a somewhat weaker generalization of Rota's NBC Theorem which, nonetheless, has some interesting connections to our work and others. To state the result we must first recall some definitions from [21]. Call $B \subseteq A(L)$ *independent* if $\bigvee B' < \bigvee B$ for every proper subset $B' \subset B$. So C is dependent if $\bigvee C' = \bigvee C$ for some $C' \subset C$. Note that these generalize the corresponding notions for a geometric lattice. The definitions of base, circuit, and broken circuit remain as before. The main result of [21] can now be stated.

Theorem 9.1 *Let L be a finite lattice. Let $\trianglelefteq$ be any total ordering of $A(L)$ such that for all ciruits C we have*

$$\bigvee C = \bigvee (C \setminus \min C). \tag{18}$$

Then for all $x \in L$ we have

$$\mu(x) = \sum_B (-1)^{|B|}$$

where the sum is over all NBC bases of x. ∎

We should note that not every lattice has a total ordering of the atoms satisfying (18) and so this result is not as strong as Theorem 1.1. On the other hand, there is an interesting relationship between a dual of this restriction and the level condition from Section 6.

Proposition 9.2 *Let L be a finite lattice. Let $\trianglelefteq$ be an induced ordering of $A(L)$ such that for all ciruits C having a unique maximal element we have*

$$\bigvee C = \bigvee (C \setminus \max C).$$

Then L satisfies the level hypothesis.

Proof. Suppose, toward a contradiction, that we have $b_0 \lhd b_1 \lhd \ldots \lhd b_k$ and $b_0 \leq \bigvee_{i\geq 1} b_i$. Then $D = \{b_0, b_1, \ldots, b_k\}$ is dependent and so contains a circuit C. Also C must have a unique maximal element b_j because D intersects each level at most once. So we have $\bigvee C = \bigvee(C \setminus b_j)$, or equivalently $b_j \leq \bigvee(C \setminus b_j)$. But this cannot happen since for some x of the inducing chain we have $b_i \leq x$ for $i < j$ but $b_j \not\leq x$. ■

In his work [1] on Tutte polynomials for hypermatroids, Christos Athanasiadis defines a *generalized lattice* to be a triple $\mathcal{L} = (L, E, f)$ where L is a finite lattice, E is a set, and $f : E \rightarrow L$ is some map. He then proves a theorem about the Möbius function of L in the case that the hypermatroid associated with $\mathcal{L}$ satifies condition (18). In the case that $E = A(L)$ and f is inclusion his result becomes a special case of Theorem 9.1.

(3) There are certain to be topological ramifications of our work. In fact Yoav Segev [22] has already proved that the order complex of any lattice is homotopy equivalent to the simplicial complex of all NBB sets B with $\bigvee B < \hat{1}$. This can be used to demonstrate a recent result of Linusson [16] that the order complex of an interval in the partition lattice under dominance is homotopy equivalent to a sphere or contractible.

An important use of NBC sets is to give a basis for the Orlik-Solomon algebra $\mathbf{A}(L)$ of a geometric lattice L [18]. If L is the intersection lattice of a complex hyperplane arrangement, then $\mathbf{A}(L)$ is isomorphic to the cohomology algebra of the complement of the arrangement. Recently there has been a lot of interest in subspace arrangements [2]. The corresponding intersection lattices are no longer geometric, but perhaps NBB sets can be used to give information about the associated cohomology algebra in this case. Recently, De Concini and Procesi [7] used algebraic geometric techniques to show that this algebra is indeed determined solely by the lattice and dimension information. This provides some hope that a combinatorial description is also possible.

References

[1] C. A. Athanasiadis, The Tutte polynomial of a hypermatroid, preprint, 1995.

[2] A. Björner, Subspace arrangements, in "Proc. 1st European Congress Math. (Paris 1992)," A. Joseph and R. Rentschler eds., Birkhäuser, Boston, MA, to appear.

[3] A. Björner and L. Lovász, Linear decision trees, subspace arrangements and Möbius functions, *J. Amer. Math. Soc.* **7** (1994), 677–706.

[4] A. Björner, L. Lovász and A. Yao, Linear decision trees: volume estimates and topological bounds, in "Proc. 24th ACM Symp. on Theory of Computing," ACM Press, New York, NY, 1992, 170–177.

[5] K. Bogart, The Möbius function of the domination lattice, unpublished manuscript, 1972.

[6] T. Brylawski, The lattice of integer partitions, *Discrete Math.* **6** (1973), 201–219.

[7] C. De Concini and C. Procesi, Wonderful models of subspace arrangements, preprint, 1995.

[8] P. H. Edelman and V. Reiner, The higher Stasheff-Tamari posets, preprint.

[9] H. Friedman and D. Tamari, Problèmes d'associativité: Une treillis finie induite par une loi demi-associative, *J. Combin. Theory* **2** (1967), 215–242.

[10] W. Geyer, On Tamari lattices, *Discrete Math.* **133** (1994), 99–122.

[11] G. Grätzer, "Lattice Theory," Freeman and Co., San Francisco, CA, 1971, pp. 17–18, problems 26-36.

[12] C. Greene, A class of lattices with Möbius function ± 1, *European J. Combin.* **9** (1988), 225–240.

[13] C. Greene, Posets of Shuffles, *J. Combin. Theory Ser. A* **47** (1988), 191–206.

[14] S. Huang and D. Tamari, Problems of associativity: A simple proof for the lattice property of systems ordered by a semi-associative law, *J. Combin. Theory Ser. A* **13** (1972), 7–13.

[15] G. Kreweras, Sur les partitions non-croisées d'un cycle, *Discrete Math.* **1** (1972), 333–350.

[16] S. Linusson, A class of lattices whose intervals are spherical or contractible, preprint.

[17] G. Markowsky, Primes, irreducibles and extremal lattices, *Order* **9** (1992), 7–13.

[18] P. Orlik and L. Solomon, Combinatorics and topology of complements of hyperplanes. *Invent. Math.* **56** (1980), 167–189.

[19] V. Reiner, Non-crossing partitions for classical reflection groups, preprint.

[20] G.-C. Rota, On the foundations of combinatorial theory I. Theory of Möbius functions, *Z. Wahrscheinlichkeitstheorie* **2** (1964), 340–368.

[21] B. E. Sagan, A generalization of Rota's NBC Theorem, *Adv. in Math.*, **111** (1995), 195–207.

[22] Y. Segev, The simplicial complex of all NBB nonspanning subsets of the set of atoms of a prelattice L is homotopic to the order complex of L, preprint, 1995.

[23] R. P. Stanley, Supersolvable lattices, *Alg. Univ.*, **2** (1972), 197–217.

[24] R. P. Stanley, "Enumerative Combinatorics, Volume 1," Wadsworth & Brooks/Cole, Pacific Grove, CA, 1986.

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

[26] T. Zaslavsky, The geometry of root systems and signed graphs, *Amer. Math. Monthly* **88** (1981), 88–105.