Partition lattice q-analogs related to q-Stirling numbers

Curtis Bennett
Department of Mathematics and Statistics
Bowling Green State University
Bowling Green, OH 43403

Kathy J. Dempsey
Department of Mathematics
Hillsdale College
Hillsdale, MI 49242

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

May 26, 1999

Key Words: set partition lattice, vector space over a finite field, q-Stirling numbers

AMS subject classification (1991): Primary 05A30; Secondary 05A18, 06A07.

Proposed running head:

Partition Lattice q-analogs

Send proofs to:

From August 1993 to August 1994:
Bruce E. Sagan
Matematiska Institutionen
Kungl. Tekniska Högskolan
Fiskartorpsv. 15A
S-100 44 Stockholm
SWEDEN
FAX: 8-723-1788

After August 1994:
Bruce E. Sagan
Department of Mathematics
Michigan State University
East Lansing, MI 48824-1027
U.S.A.
FAX: 517-336-1562

Abstract

We construct a family of partially ordered sets (posets) that are q-analogs of the set partition lattice. They are different from the q-analogs proposed by Dowling [Dow 73]. One of the important features of these posets is that their Whitney numbers of the first and second kind are just the q-Stirling numbers of the first and second kind, respectively. One member of this family [Dem 92] can be constructed using an interpretation of Milne [Mil 82] for $S[n, k]$ as sequences of lines in a vector space over the Galois field F_q. Another member is constructed so as to mirror the partial order in the subspace lattice.

1 Introduction and definitions

The theory of q-analogs has a long and venerable history, going back to Gauss who studied the q-binomial coefficients. Recently there has been considerable interest in q-Stirling numbers. The q-Stirling numbers of the first kind, $s[n,k]$, were introduced in Gould's paper [Gou 61]. Then Gessel [Ges 82] gave them an interpretation as the generating functions for an inversion statistic. The q-Stirling numbers of the second kind, $S[n,k]$, were introduced by Carlitz [Car 33, Car 48], and were also studied by Gould [Gou 61] and Milne [Mil 78]. Later, Milne [Mil 82] showed that these numbers could be viewed in terms of inversions on partitions, and Sagan [Sag 91] showed that there was also a version of the major index. Also, Milne [Mil 82] and Wachs and White [W-W 91] proved that the q-Stirling numbers of the second kind count restricted growth functions using various statistics. Finally, Milne [Mil 82] gave a vector space interpretation of these numbers.

Consider the lattice of partitions of $\{1,2,\ldots,n\}$ ordered by refinement, Π_n. This lattice has the ordinary Stirling numbers of the first and second kinds, $s(n,k)$ and $S(n,k)$, as its Whitney numbers of the first and second kinds, respectively. The purpose of this paper is to construct a q-analog of Π_n that has $s[n,k]$ and $S[n,k]$ as its Whitney numbers. (Dowling [Dow 73] has also constructed a q-partition lattice, but it does not have this property.) In fact we will construct a whole family of such q-analogs. One in particular will use Milne's interpretation of the $S[n,k]$ as sequences of lines in a vector space over a finite field. Another will reflect a connection with the lattice of subspaces. First, however, we will need some definitions and notation.

Let $\mathbf{Z}$ and $\mathbf{N}$ stand for the integers and non-negative integers, respectively. Consider $\hat{n} = \{1,\ldots,n\}$. The set of all permutations of $\hat{n}$ that can be written as a product of k disjoint cycles will be denoted $c(\hat{n},k)$. The *(ordinary) Stirling numbers of the first kind*, $s(n,k)$, are defined by

$$s(n,k) = (-1)^{n-k}|c(\hat{n},k)|$$

where $|\cdot|$ denotes cardinality. It is easy to see that these numbers satisfy the recursion

$$s(n,k) = \begin{cases} s(n-1,k-1) - (n-1)\, s(n-1,k) & \text{for } n \geq 1 \\ \delta_{0,k} & \text{for } n = 0. \end{cases} \tag{1}$$

where $\delta_{0,k}$ is the Kronecker delta, with k and n running over $\mathbf{Z}$ and $\mathbf{N}$, respectively.

If $n \in \mathbf{N}$ then it's standard q-analog is

$$[n] = 1 + q + q^2 + \cdots + q^{n-1}.$$

So we can define the *q-Stirling numbers of the first kind*, $s[n,k]$, using a q-analog of equation (1):

$$s[n,k] = \begin{cases} s[n-1,k-1] - [n-1]\, s[n-1,k] & \text{for } n \geq 1 \\ \delta_{0,k} & \text{for } n = 0. \end{cases} \tag{2}$$

Next we will look at Stirling numbers of the second kind. The set of all partitions of $\hat{n}$ into disjoint subsets $B_1, B_2, \ldots, B_k$ (sometimes called parts or blocks) will be denoted $S(\hat{n},k)$. The *(ordinary) Stirling numbers of the second kind* are

$$S(n,k) = |S(\hat{n},k)|.$$

The recurrence for these numbers is

$$S(n,k) = \begin{cases} S(n-1,k-1) + k\, S(n-1,k) & \text{for } n \geq 1 \\ \delta_{0,k} & \text{for } n = 0. \end{cases} \tag{3}$$

with the usual conventions for k and n. A q-analog of this equation gives the *q-Stirling numbers of the second kind*

$$S[n,k] = \begin{cases} S[n-1,k-1] + [k]\, S[n-1,k] & \text{for } n \geq 1 \\ \delta_{0,k} & \text{for } n = 0. \end{cases} \tag{4}$$

There are many different statistics that yield the q-Stirling numbers [Ges 82, Mil 82, Sag 91, W-W 91]. The one that will interest us counts non-inversions. If $\pi \in S(\hat{n},k)$, then we write

$$\pi = B_1/B_2/\cdots/B_k$$

with the convention that

$$1 = \min B_1 < \min B_2 \cdots < \min B_k.$$

We will say that a partition written in this manner is in *standard form*. Given such a partition π we will let $m_i = \min B_i$ and denote by b_i any element of $B_i \setminus \{m_i\}$. It should also be noted that we always refer to elements of π with small letters, and to blocks of π with capital letters.

Define a *non-inversion* of π to be a pair (m_i, b_j) such that $m_i < b_j$ and $i < j$. The *non-inversion set of* π is

$$\mathrm{Nin}\,\pi = \{(m_i, b_j) \mid (m_i, b_j) \text{ is a non-inversion in } \pi\}$$

and we let

$$\mathrm{nin}\,\pi = |\,\mathrm{Nin}\,\pi|.$$

If (a, b) is a non-inversion in π we will often say that b *causes a non-inversion* in π. Similarly if there are l different elements a such that (a, b) is a non-inversion then we will say that b *causes l non-inversions.* Note that

$$\operatorname{nin} \pi = \sum_{j=1}^{k} (|B_j| - 1)(j - 1) \tag{5}$$

since each b_j causes a non-inversion with all m_i such that $i < j$.

As an example, consider

$$\pi = 139/2467/58. \tag{6}$$

Then

$$\operatorname{Nin} \pi = \{(1, 4), (1, 6), (1, 7), (1, 8), (2, 8)\}$$

and $\operatorname{nin} \pi = 5$. Also we can say 8 causes two non-inversions and 7 causes one non-inversion.

We will now show that $S[n, k]$ is the generating function for nin.

Theorem 1.1 *We have*

$$S[n, k] = \sum_{\pi \in S(\hat{n}, k)} q^{\operatorname{nin} \pi}$$

Proof. It suffices to check the boundary condition and recursion of (4). The former is easy. For the latter, note that

$$\begin{aligned} \sum_{\pi \in S(\hat{n}, k)} q^{\operatorname{nin} \pi} &= \sum_{B_k = \{n\}} q^{\operatorname{nin} \pi} + \sum_{B_k \neq \{n\}} q^{\operatorname{nin} \pi} \\ &= \sum_{\pi' \in S(\widehat{n-1}, k-1)} q^{\operatorname{nin} \pi'} + (1 + q + \cdots + q^{k-1}) \sum_{\pi' \in S(\widehat{n-1}, k)} q^{\operatorname{nin} \pi'} \end{aligned}$$

where the $1 + q + \cdots + q^{k-1}$ comes from the non-inversions caused by the element n. Hence by induction

$$\begin{aligned} \sum_{\pi \in S(\hat{n}, k)} q^{\operatorname{nin} \pi} &= \sum_{\pi' \in S(\widehat{n-1}, k-1)} q^{\operatorname{nin} \pi'} + [k] \sum_{\pi' \in S(\widehat{n-1}, k)} q^{\operatorname{nin} \pi'} \\ &= S[n-1, k-1] + [k] S[n-1, k] \\ &= S[n, k]. \quad \blacksquare \end{aligned}$$

We will also wish to compare two non-inversion sets. If $\sigma, \pi \in \Pi_n$, then the set of *new non-inversions of* $\sigma \backslash \pi$ is defined to be the set-theoretic difference

$$\operatorname{Nin} \sigma \backslash \pi = (\operatorname{Nin} \sigma) \setminus (\operatorname{Nin} \pi).$$

In addition, let $\operatorname{nin}\sigma\backslash\pi = |\operatorname{Nin}\sigma \setminus \pi|$. To illustrate, if

$$\sigma = 139/27/46/58 \tag{7}$$

and π is as in (6), then

$$\operatorname{Nin}\sigma \setminus \pi = \{(2,6),(4,8)\}.$$

with $\operatorname{nin}\sigma \setminus \pi = 2$.

Finally, we need some definitions concerning posets (partially ordered sets). Consider a finite poset P with unique minimum $\hat{0}$ and maximum $\hat{1}$. Define the *Möbius function* $\mu : P \to \mathbf{Z}$ by

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

The partitions of $\hat{n}$ form a poset, Π_n, when ordered by refinement. The Möbius function of Π_n is well known. The reader can consult Stanley's text [Sta 86] for a proof.

Theorem 1.2 *If $\pi = B_1/B_2/\ldots/B_k$ is an element of Π_n, then*

$$\mu(\pi) = (-1)^{n-k}\prod_{i=1}^{k}(|B_i|-1)! \ \blacksquare$$

A poset P is *ranked* if, for every element $x \in P$, all maximal chains from $\hat{0}$ to x are of the same length. The *rank of x* is this common length. If a poset is ranked, then we can also consider its Whitney numbers of the first and second kinds. The k^{th} *Whitney number of the first kind* is

$$w_k(P) = \sum_{\operatorname{rk} x = k}\mu(x),$$

so it is the sum of all the Möbius function values at rank k. The k^{th} *Whitney number of the second kind* is

$$W_k(P) = \sum_{\operatorname{rk} x = k} 1,$$

so it is the number of elements at rank k. The connection between the Whitney numbers and the Stirling numbers is provided by the following theorem. Again, see [Sta 86] for details.

Theorem 1.3 *The Whitney numbers for* Π_n *are*

$$w_k(\Pi_n) = s(n, n-k) \quad \textit{and} \quad W_k(\Pi_n) = S(n, n-k). \blacksquare$$

Our aim in the rest of this paper is to describe a family of posets, $\mathcal{P}_n(q)$, each of which can claim to be a q-analog of Π_n. The basis of this claim is that for each $P \in \mathcal{P}_n(q)$, a q-analog of Theorem 1.3 will be true with the Stirling numbers replaced by q-Stirling numbers. Furthermore, a q-version of Theorem 1.2 will be true. Later we will see what replaces the factorials therein. First, however, we will examine two particular posets in $\mathcal{P}_n(q)$ that are of special interest.

2 The poset $\Pi_n(q)$

For the purposes of this section and the next, let q be a prime power. Milne [Mil 78] gave a vector space interpretation to the $S[n,k]$ as follows. Let $\mathbf{V}$ be a vector space over the Galois field F_q. Given any set of vectors $S \subseteq \mathbf{V}$, we let $\langle S \rangle$ denote the linear span of S. Now any sequence of lines (one dimensional subspaces of $\mathbf{V}$) $l_1, l_2, \ldots, l_n$ determines a *flag*

$$\mathbf{V}_0 \subset \mathbf{V}_1 \subset \ldots \subset \mathbf{V}_k \tag{8}$$

obtained by deleting the repeated subspaces in

$$\langle 0 \rangle \subseteq \langle l_1 \rangle \subseteq \ldots \subseteq \langle l_1, l_2, \ldots, l_n \rangle.$$

The sequence (8) is called a *complete flag of dimension* k because $\dim \mathbf{V}_i = i$ for $1 \leq i \leq k$. Now fix a complete flag of dimension k, and let $\{v_1, \ldots, v_k\}$ be a basis for $\mathbf{V}_k$ such that $\{v_1, \ldots, v_i\}$ is a basis for $\mathbf{V}_i$ for all i. Then the group of upper unitriangular matrices in this basis acts on the sequences of n lines generating the given flag of dimension k. Milne showed that $S[n,k]$ is the number of orbits under this action, regardless of the initial choice of a flag.

We would like a more concrete description of this situation. Therefore, we will pick a particular representative for each orbit. By a common abuse of notation, we will represent a line l by a non-zero vector lying along l. We may assume that the vector representing a line is of the form $(*, \ldots, *, 1, 0, \ldots, 0)$ where each $*$ is an element of F_q. We will say that the rightmost non-zero entry of this vector is its *leading 1*.

Since we can pick any complete flag of dimension k, we will do so in the simplest possible way. Let $\epsilon_1, \ldots, \epsilon_n$ be the standard basis vectors for F_q^n. Then define

$$\mathbf{W}_i = \langle \epsilon_1, \ldots, \epsilon_i \rangle.$$

These will be the subspaces in our complete flag. To account for the orbits, we will put each sequence of lines generating $\mathbf{W}_0 \subset \mathbf{W}_1 \subset \ldots \subset \mathbf{W}_k$ in a canonical form. So let $\ell = l_1, \ldots, l_n$ be such a sequence. We will say that ℓ is *standard* if for each i with $l_i \notin \langle l_1, \ldots, l_{i-1} \rangle$ we have $l_i = \epsilon_j$ where $\langle l_1, \ldots, l_i \rangle = \mathbf{W}_j$. Further, we say the *sequence has dimension* k if the flag generated is of dimension k.

For example, if $n = 9, k = 4$ and $q = 3$, then one possible sequence of lines is

$$\ell = \begin{array}{c} (1,0,0,0,0,0,0,0,0) \\ (0,1,0,0,0,0,0,0,0) \\ (1,0,0,0,0,0,0,0,0) \\ (0,0,1,0,0,0,0,0,0) \\ (0,0,0,1,0,0,0,0,0) \\ (2,1,1,0,0,0,0,0,0) \\ (2,1,0,0,0,0,0,0,0) \\ (1,0,2,1,0,0,0,0,0) \\ (1,0,0,0,0,0,0,0,0) \end{array} \tag{9}$$

Note that the first, second, fourth and fifth lines must be standard basis vectors since, in each case, they are the first to venture into a larger subspace than their predecessors. Notice, too, that our definition makes sense for all positive integers, q, not just prime powers. Merely replace elements of F_q with the residue class $\{0, 1, \ldots, q-1\}$ of integers modulo q. As long as we do not try to do any division, there will be no problems. Hence all our enumerative results will be true in this larger setting.

Let $\Pi_n(q)$ denote the set of all standard sequences of n lines and consider $\ell = l_1, \ldots, l_n \in \Pi_n(q)$. Define a mapping $T : \Pi_n(q) \to \Pi_n$ by $T(\ell) = B_1/\ldots/B_k$ where

$$i \in B_j \Leftrightarrow \text{the leading 1 of } l_i \text{ is in position } j.$$

This establishes a surjective correspondence between sequences of lines, ℓ, and partitions, π. We call T the *type map*, and say that ℓ has *type* π. Because of this correspondence, we will say that l_i is *in the* j^{th} *part* whenever its leading 1 is in position j. In other words, l_i is in the j^{th} part whenever $l_i \in \mathbf{W}_j \setminus \mathbf{W}_{j-1}$. By way of illustration, if ℓ is the sequence (9) then

$$T(\ell) = 139/27/46/58.$$

We can now state the connection between line sequences and the non-inversion statistic.

Proposition 2.1 *Given $\pi \in \Pi_n$, the number of $\ell \in \Pi_n(q)$ of type π is $q^{\operatorname{nin} \pi}$.*

Proof. Suppose $T(\ell) = \pi = B_1/\ldots/B_k$ and consider i, $1 \leq i \leq n$. If i is the minimum of a block of π, then l_i must be a standard basis vector. So there is no choice for these lines in ℓ.

If $i \in B_j \setminus \{m_j\}$ then i causes $j-1$ non-inversions in π. Also, l_i has $j-1$ entries which are arbitrary elements of F_q, so there are q^{j-1} possibilities for l_i. Thus, the total number of possible preimages of π is

$$\prod_{j\geq 1} q^{(|B_j|-1)(j-1)} = q^{\sum_{j\geq 1}(|B_j|-1)(j-1)} = q^{\operatorname{nin}\pi}$$

by equation (5). ■

Given a line $l = (a_1, \ldots, a_{k-1}, 1, 0, \ldots, 0)$ define the *shift left operator*, L, by

$$Ll = (a_2, \ldots, a_{k-1}, 1, 0, \ldots, 0)$$

so that the 1 now appears in the $(k-1)^{\text{st}}$ position and there is one extra 0 at the end. *Shifting left* m, $L^m l$, is just the result of shifting left m times in succession. If l is the eighth line of (9) then $Ll = (0, 2, 1, 0, 0, 0, 0, 0, 0)$ and $L^2 l = (2, 1, 0, 0, 0, 0, 0, 0, 0)$.

We will now define a partial order on standard sequences of n lines to turn $\Pi_n(q)$ into a poset. We wish to make the covering relation in $\Pi_n(q)$ agree with the the covering relation in Π_n. Given $\ell = l_1, \ldots, l_n$ and $\ell' = l'_1, \ldots, l'_n$, we write $\ell' \prec \ell$ if ℓ' is covered by ℓ. So we want

$$\ell' \prec \ell \Rightarrow T(\ell') \prec T(\ell).$$

If σ is a partition, then we make a covering partition by combining two parts, say the s^{th} and t^{th} parts where $s < t$. Hence for sequences of lines we will move lines in the t^{th} into the s^{th} part. Note also that the indices on blocks after the t^{th} in σ are decreased by one, whereas the indices of those blocks before the t^{th} remain the same. This will be mirrored as well in our definition of the covering relation.

Now suppose $\ell = l_1, \ldots, l_n$ and $\ell' = l'_1, \ldots, l'_n$ are standard. Then we let $\ell' \prec \ell$ if there exist s and t, $1 \leq s < t \leq n$, such that the following four conditions hold.

1. If l'_m is the first line in the t^{th} part then $l_m = (a_1, \ldots, a_{s-1}, 1, 0, \ldots, 0)$ where the $a_i \in F_q$ are arbitrary.

2. For each l'_i in the t^{th} part with $i \neq m$, we have $l_i = L^{t-s} l'_i$.

3. For each l'_i in a part after the t^{th}, we have $l_i = L l'_i$.

4. For each l'_i in a part before the t^{th}, $l_i = l'_i$.

Continuing our running example,

$$\begin{array}{l}(1,0,0,0,0,0,0,0,0)\\(0,1,0,0,0,0,0,0,0)\\(1,0,0,0,0,0,0,0,0)\\(0,0,1,0,0,0,0,0,0)\\(0,0,0,1,0,0,0,0,0)\\(2,1,1,0,0,0,0,0,0)\\(2,1,0,0,0,0,0,0,0)\\(1,0,2,1,0,0,0,0,0)\\(1,0,0,0,0,0,0,0,0)\end{array} \quad \stackrel{T}{\longrightarrow} \quad 139/27/46/58$$

is covered by

$$\begin{array}{l}(1,0,0,0,0,0,0,0,0)\\(0,1,0,0,0,0,0,0,0)\\(1,0,0,0,0,0,0,0,0)\\(*,1,0,0,0,0,0,0,0)\\(0,0,1,0,0,0,0,0,0)\\(1,1,0,0,0,0,0,0,0)\\(2,1,0,0,0,0,0,0,0)\\(0,2,1,0,0,0,0,0,0)\\(1,0,0,0,0,0,0,0,0)\end{array} \quad \stackrel{T}{\longrightarrow} \quad 139/2467/58$$

where the * could be any element of F_3.

The proof of the next result follows directly from the preceding definition and so is omitted.

Proposition 2.2 *The map $T : \Pi_n(q) \to \Pi_n$ is order-preserving.* ■

Finally, we can use new non-inversions to count the number of lines below a given element of $\Pi_n(q)$.

Proposition 2.3 *Fix $\ell \in \Pi_n(q)$ of type π and $\sigma \leq \pi$. Then the number of $\ell' \leq \ell$ with ℓ' of type σ is $q^{\operatorname{nin} \sigma \setminus \pi}$.*

Proof. Suppose $\pi = B_1/B_2/\ldots$ and $i \in B_j$. If i is not a minimum in $\sigma = C_1/C_2/\ldots$ and $i \in C_m$, then i causes $m-j$ of the non-inversions in $\operatorname{Nin} \sigma \setminus \pi$. Furthermore, these are the only i, $1 \leq i \leq n$, that cause new non-inversions.

Now suppose that $\ell' \leq \ell$ is of type σ, and consider the line l_i in ℓ. By definition, l_i completely determines the corresponding l_i' in ℓ' unless i is not a minimum in σ. In that case, l_i' will have $m-j$ entries that can be chosen arbitrarily if it moves from the j^{th} part of π to the m^{th} part of σ. Thus the number of choices for l_i' is q^{m-j}. But this $m-j$ coincides with the number of new non-inversions computed in the previous paragraph. So the total number of lines will be $q^{\operatorname{nin} \sigma \setminus \pi}$. ■

3 The poset $P_n(q)$

We define a poset $P_n(q)$ as follows. Let

$$\widehat{F}_q = F_q \cup \{\infty\}.$$

Extend the addition and multiplication in F_q to $\widehat{F}_q$ by letting

$$\begin{array}{l} \infty + x = x + \infty = \infty \\ x \cdot \infty = \infty \cdot x = 0 \end{array} \tag{10}$$

for all $x \in \widehat{F}_q$. Let $P_n(q)$ consist of the $(k-1)\times(n-1)$ matrices over $\widehat{F}_q$, $1 \le k \le n$, in the following version of row-echelon form.

1. Each row has at least one finite entry, and the first finite entry in each row is a 1 called a *pivot*.
2. All entries above a pivot are 0, and all entries below a pivot are ∞.
3. All entries below an ∞ are ∞.

We will find it convenient to index the rows with $2, \ldots, k$ and the columns with $2, \ldots, n$. Again, as long as division is not used we can replace F_q with a complete set of residues modulo $q \in \mathbf{N}$.

By way of illustration, if $n = 9$, $k = 4$ and $q = 3$ then one such matrix is

$$M = \begin{pmatrix} 1 & \infty & 0 & 0 & 1 & 1 & 0 & \infty \\ \infty & \infty & 1 & 0 & 0 & \infty & 1 & \infty \\ \infty & \infty & \infty & 1 & \infty & \infty & 2 & \infty \end{pmatrix}. \tag{11}$$

The unique two in M is $m_{4,8} = 2$.

We now define a *type map* $T : P_n(q) \to \Pi_n$ by letting $T(M) = B_1/\ldots/B_k$ where

$$\begin{array}{lcl} j \in B_1 & \Leftrightarrow & j = 1 \text{ or } m_{i,j} = \infty \text{ for all } i \\ j \in B_i & \Leftrightarrow & i \text{ is maximal such that } m_{i,j} \neq \infty \end{array}$$

If M is as in (11), then

$$T(M) = 139/27/46/58.$$

If $T(M) = \pi$ then we will say that M *has type* π.

We have an analog of Proposition 2.1 for $P_n(q)$.

Proposition 3.1 *Given $\pi \in \Pi_n$, then number of $M \in P_n(q)$ of type π is $q^{\mathrm{nin}\,\pi}$.*

Proof. Suppose $T(M) = \pi = B_1/B_2/\ldots$ and consider j, $1 \le j \le n$. If j is the minimum of a block of π, then the j^{th} column of M contains a pivot 1. So all the entries of this column are fixed as a 0, 1 or ∞. Thus there is no choice for these columns in M.

If $j \in B_i \setminus \{m_i\}$ then j causes $i-1$ non-inversions in π. Also, in the j^{th} column of M, the first $i-1$ entries are arbitrary (and the rest fixed), so there are q^{i-1} possibilities for that column. Thus, the total number of possible preimages of π is

$$\prod_{i\ge 1} q^{(|B_i|-1)(i-1)} = q^{\sum_{i\ge 1}(|B_i|-1)(i-1)} = q^{\operatorname{nin}\pi}$$

by equation (5). ■

For brevity, we will write a $(k-1)\times(n-1)$ matrix M using row vectors. That is,

$$M = \begin{pmatrix} v_2 \\ \vdots \\ v_k \end{pmatrix},$$

where the v_i are vectors over $\widehat{F}_q$. Vector addition and scalar multiplication are defined componentwise using (10) with one exception. If $T(M) = B_1/\ldots/B_k$, then we let

$$\infty \cdot v_t = (a_{t,2}, \ldots, a_{t,n})$$

where

$$a_{t,j} = \begin{cases} 0 & \text{if } j \notin B_t \\ \infty & \text{if } j \in B_t \end{cases} \tag{12}$$

We are now ready to define the covers of the poset $P_n(q)$. Let

$$M' = \begin{pmatrix} v'_2 \\ \vdots \\ v'_{k+1} \end{pmatrix}. \tag{13}$$

Then M covers M' if and only if there exist $1 \le s < t \le k+1$ such that

$$M = \begin{pmatrix} v'_2 + x_2 v'_t \\ \vdots \\ v'_{t-1} + x_{t-1} v'_t \\ v'_{t+1} \\ \vdots \\ v'_{k+1} \end{pmatrix}, \tag{14}$$

where $x_i \neq \infty$ for $i \leq s$, and $x_i = \infty$ for $i > s$.

Continuing our previous example, consider

$$\begin{pmatrix} v_2' \\ v_3' \\ v_4' \end{pmatrix} = \begin{pmatrix} 1 & \infty & 0 & 0 & 1 & 1 & 0 & \infty \\ \infty & \infty & 1 & 0 & 0 & \infty & 1 & \infty \\ \infty & \infty & \infty & 1 & \infty & \infty & 2 & \infty \end{pmatrix} \xrightarrow{T} 139/27/46/58.$$

Letting $s = 2$ and $t = 4$, we see that this matrix is covered by

$$\begin{pmatrix} v_2' + 2v_4' \\ v_3' + \infty v_4' \end{pmatrix} = \begin{pmatrix} 1 & \infty & 0 & 2 & 1 & 1 & 1 & \infty \\ \infty & \infty & 1 & \infty & 0 & \infty & \infty & \infty \end{pmatrix} \xrightarrow{T} 139/2578/46$$

where all addition and multiplication is done in $\hat{F}_3$. The fact that a covering matrix is sent by T to a covering partition is no accident, as the next proposition shows.

Proposition 3.2 *If M' and M are as in (13) and (14), then M is in row-echelon form. Furthermore, the map $T : P_n(q) \to \Pi_n$ is order-preserving.*

Proof. We will simultaneously prove that M is well-defined, and that if $T(M') = B_1/\ldots/B_{k+1}$, then

$$T(M) = B_1/\ldots/B_{s-1}/B_s \cup B_t/B_{s+1}/\ldots/B_{t-1}/B_{t+1}/\ldots/B_{k+1}. \tag{15}$$

This will give us the order-preserving property of T as well.

First note that for $i \leq s$ all entries of $x_i v_t'$ are finite because of the multiplication rules in (10) and the fact that $x_i \neq \infty$. So v_i' and $v_i' + x_i v_t'$ have ∞'s in exactly the same spots. Thus the upper portion of M is well-defined. Applying the type map to the matrix consisting of rows 2 through s of M yields the blocks $B_2, \ldots, B_{s-1}$ which are the same as for M'.

Now $\infty v_t'$ has ∞ in all positions j such that $j \in B_t$ and zeros elsewhere. Thus, for $s < i < t$, the vector $v_i' + \infty v_t'$ will have extra ∞'s in the positions of $j \in B_t$, and the same elements as v_i' in the other positions. Hence rows $2, \ldots, t-1$ of M are in row-echelon form and the type map applied to these rows gives the second through $(t-2)^{\text{nd}}$ blocks of (15).

Finally, v_i' is invariant for $i \geq t+1$. This completes the verification that M is well-defined and that $T(M)$ also has blocks B_1 and $B_{t+1}, \ldots, B_{k+1}$. ∎

We will also need an analog of Proposition 2.3 in this setting.

Proposition 3.3 *Fix $M \in P_n(q)$ of type π and $\sigma \leq \pi$. Then the number of $M' \leq M$ with M' of type σ is $q^{\operatorname{nin} \sigma \backslash \pi}$.*

Proof. We will first reduce to the case of a cover. Let the elements which are minima in σ but not in π be $s_1 > s_2 > \dots$ and construct a sequence of partitions

$$\pi = \pi_0 \succ \pi_1 \succ \pi_2 \succ \dots \succ \sigma$$

where π_l is obtained from π_{l-1} by making s_l a new minimum. Because of our choice of covers, we have

$$\operatorname{Nin} \pi \setminus \sigma = \biguplus_{l \geq 1} \operatorname{Nin} \pi_l \setminus \pi_{l-1}$$

where $\uplus$ denotes disjoint union. Now for any $M' < M$ with $T(M') = \sigma$ there is a sequence of matrices in $P_n(q)$

$$M = M_0 \succ M_1 \succ M_2 \succ \dots \succ M'$$

such that $T(M_l) = \pi_l$ for all l. Thus it suffices to consider covers.

So suppose that $M' \prec M$ and that blocks B_s and B_t, $s < t$, are merged in $\sigma = B_1/\dots/B_{k+1}$ to form π. Note that the entries in v_i' for $i > t$ are not changed in passing from M' to M. So we need only consider $i \leq t$.

Now consider column j of M'. If j is a minimum of σ then it causes no new non-inversions in $\operatorname{Nin} \sigma \setminus \pi$. Such elements correspond to pivot columns of M' where every element is determined. Thus there is only $q^0 = 1$ possibility for such columns, as desired.

If j is in a block to the left of B_t, then it also causes no new non-inversions. In this case the j^{th} element of row v_t' in M' is equal to ∞. Also, since $j \notin B_t$, adding multiples of v_t does not affect the j^{th} column by (10) and (12). Thus the entries of this column in M completely determine the entries of this column in M'. So, again, there is only one choice.

If $j \in B_t \backslash \{m_t\}$, then j causes $t-s$ new non-inversions. We claim that the entries $m'_{s+1,j}, m'_{s+2,j}, \dots, m'_{t,j}$ can be chosen arbitrarily from F_q, and once this is done all the other entries of the j^{th} column of M' are determined uniquely. This will give the necessary q^{t-s} possibilities. Since for $s < i < t$ we have $\infty = m_{i,j} = m'_{i,j} + \infty$, any finite entry will do for the $m'_{i,j}$ in these rows. This accounts for q^{t-s-1} choices. For $i \leq s$, note that the scalar $x_i \neq \infty$ in (14) is uniquely determined by the entries in column m_t of M and the fact that v_t' has its pivot one in that column. So, given $m'_{t,j}, m_{i,j} \neq \infty$ there is a unique solution to

$$m_{i,j} = m'_{i,j} + x_i m'_{t,j}. \tag{16}$$

Thus, any one of the q choices for $m'_{t,j}$ completely determines the rest of column j. Hence, the total is $q^{t-s-1} \cdot q = q^{t-s}$ possibilities.

If j is a non-minimum in a block to the right of B_t, then it causes one new non-inversion. For $s < i < t$ we must have $m'_{i,j} = m_{i,j}$ by (12), so these entries are uniquely determined. Also, once we have picked $m'_{t,j}$, the entries $m'_{i,j}$ for $i \leq s$ are fixed by (16). This gives a total of q^1 choices, and we have checked that the proposition holds in every case. Note that the solution of (16) for $m'_{i,j}$, given all the other quantities in the equation, can be done without division. Thus this result still holds when q is not a prime power. ■

We will now show that the subspace lattice of dimension $n-1$ over F_q, $L_{n-1}(q)$, can be embedded in a natural way in $P_n(q)$. Given a basis, each element of the subspace lattice corresponds to a matrix in standard row-echelon form. So, replacing each 0 to the left of a pivotal 1 with ∞, we can obtain an element of $P_n(q)$. This defines a map $\psi : L_{n-1}(q) \to P_n(q)$. It is easy to see that if $\mathbf{V} \in L_{n-1}(q)$ then $\psi(\mathbf{V})$ has the correct form for a row-echelon matrix in $P_n(q)$. We claim that the image of ψ is a subposet of $P_n(q)$ isomorphic to $L_{n-1}(q)$. This will follow from the next proposition and the fact that the subspace lattice is self-dual.

Proposition 3.4 *The map $\psi : L_{n-1}(q) \to P_n(q)$ is an order-reversing injection. Also, ψ^{-1} is order-reversing.*

Proof. The fact that ψ is injective follows immediately from its definition. The proofs that ψ and ψ^{-1} are order-reversing are similar, so we will only do the former.

It suffices to show that if $\mathbf{V} \prec \mathbf{W}$ then $\psi(\mathbf{V}) \succ \psi(\mathbf{W})$. But if

$$\mathbf{W} = \begin{pmatrix} w_1 \\ \vdots \\ w_k \end{pmatrix}$$

in row echelon form, then it is possible to find scalars $x_1, \ldots, x_{t-1} \in F_q$ such that

$$\mathbf{V} = \begin{pmatrix} w_1 + x_1 w_t \\ \vdots \\ w_{t-1} + x_{t-1} w_t \\ w_{t+1} \\ \vdots \\ w_k \end{pmatrix}$$

for some t. This corresponds to the case $s = t - 1$ in the definition of the covers for $P_n(q)$. Thus $\psi(\mathbf{V}) \succ \psi(\mathbf{W})$. ■

4 Properties of Π_n

We will now investigate some properties of the partition lattice that will be important for our q-analogs. Let $\pi = B_1/B_2/\ldots/B_k$ be a partition of $\hat{n}$. Recall the convention that $m_i = \min B_i$ and $b_i \in B_i \setminus \{m_i\}$.

We will need to refine the set of non-inversions. So define the *non-inversions caused by a block*, B_j, by

$$\operatorname{Nin}_\pi B_j = \{(m_i, b_j) \mid m_i < b_j \text{ and } i < j\}.$$

Similarly, define the *non-inversions caused by an element*, $b_j \in \hat{n}$, by

$$\operatorname{Nin}_\pi b_j = \{(m_i, b_j) \mid m_i < b_j \text{ and } i < j\}.$$

For example, if $\pi = 139/2467/58 = B_1/B_2/B_3$, then

$$\operatorname{Nin} B_2 = \{(1,4),(1,6),(1,7)\}$$

and

$$\operatorname{Nin} 8 = \{(1,8),(2,8)\}.$$

Clearly

$$\operatorname{Nin} \pi = \biguplus_{B \in \pi} \operatorname{Nin}_\pi B$$

and

$$\operatorname{Nin}_\pi B = \biguplus_{b \in B \setminus \{m\}} \operatorname{nin}_\pi b.$$

As before, "nin" will be used to denote the number of elements in each set. Also the π will be dropped when doing so causes no confusion.

The *interference set of* π is

$$\operatorname{Int} \pi = \{(b_i, b_j) \mid b_i < m_j \text{ and } i < j\}$$

with cardinality

$$\operatorname{int} \pi = |\operatorname{Int} \pi|.$$

Using the same partition as before

$$\operatorname{Int} \pi = \{(3,8),(4,8)\}.$$

If $(b_i, b_j) \in \operatorname{Int} \pi$, then we say (as usual) b_j *causes interference in* π. Note that

$$\operatorname{int} \pi = \sum_{i=2}^{k} (|B_i| - 1)(\min B_i - i)$$

since each element that is not the first in its block causes interference with each element smaller than the minimum of its block, except those that are themselves first in a block. We will also refine interference by defining

$$\operatorname{Int}_\pi B_j = \{(b_i, b_j) \mid b_i < m_j \text{ and } i < j\}$$

and

$$\operatorname{Int}_\pi b_j = \{(b_i, b_j) \mid b_i < m_j \text{ and } i < j\}.$$

As previously, "int" will denote the number of elements in each set and the π will be dropped when doing so causes no confusion.

Intuitively, (b_i, b_j) is an interference pair if splitting block B_i up into B_{i_1} and B_{i_2} with $b_i = \min B_{i_2}$ will make (b_i, b_j) a new non-inversion. This is the essential ingredient in the following lemma.

Lemma 4.1 *Let π be a partition and let B be a block of π. If $b \notin B$, then*

$$\operatorname{int}_\pi b = \operatorname{int}_\sigma b + \operatorname{nin}_{\sigma\backslash\pi} b$$

for any refinement, σ, of π which splits B into two parts and keeps all other parts the same.

Proof. Suppose the pair (a, b) causes interference in π where $b \notin B$. Let B' be the block of b in π. Then B' is still b's block in σ. Furthermore, since $a < m = \min B'$, we have that the block containing a still precedes B' in σ. Now there are two possibilities for a. If a becomes a minimum of a block in σ, then (a, b) is a new non-inversion in $\sigma\backslash\pi$. Otherwise, (a, b) still causes interference in σ. Thus we have shown

$$\operatorname{int}_\pi b \leq \operatorname{int}_\sigma b + \operatorname{nin}_{\sigma\backslash\pi} b.$$

For the reverse inequality, suppose $b \in \operatorname{Int}_\sigma b \uplus \operatorname{Nin}_{\sigma\backslash\pi} b$. If (a, b) causes interference in σ, then it also causes interference in π. If (a, b) is a new non-inversion, then a must be the larger minimum in the two blocks merged to form B. Thus these pairs also cause interference in π. This concludes the proof of the lemma. ∎

In order to obtain q-analogs of the factorials appearing in the Möbius function for the partition lattice, we will need a rather strange definition. Given any set of integers, $B = \{a_1 < a_2 < \ldots < a_n\}$ define

$$B! = q^{a_1}(q^{a_1} + q^{a_2}) \cdots (q^{a_1} + q^{a_2} + \cdots + q^{a_n}).$$

Thus $\{0, 1, \ldots, n-1\}!$ is the usual q-analog of $n!$. Also, let

$$B - b = \{a_1 - b, a_2 - b, \ldots, a_{n-1} - b\}.$$

It should be pointed out that $B-b$ has only $n-1$ elements and that $(B-b)!$ is a q-analog of $(n-1)!$. Note that we can consider $(B-b)!$ to be a product with a factor corresponding to each element of $B\setminus\min B$ where the *factor corresponding to* a_j, $j\geq 2$, is

$$q^{a_1-b}+q^{a_2-b}+\cdots+q^{a_{j-1}-b}.$$

We will need a lemma telling us how this factorial acts under refinement.

Lemma 4.2 *Let $B=\{a_1<a_2<\ldots<a_n\}$ be a set of positive integers, then*

$$(B-a_1)!=\sum(C-a_1)!(D-a_1)!$$

where the sum is taken over all partitions of the set B into two parts $B=C\uplus D$ with $a_1\in C$ and $a_2\in D$.

Proof. Assume that these products are multiplied out to give monomials but terms with equal powers are not combined. We will first show that there are the same number of terms on each side. Then, since the sums are finite, it will be sufficient to prove that each term of $(B-a_1)!$ appears as a unique term in $(C-a_1)!(D-a_1)!$ for some uniquely chosen C and D.

Now $(B-a_1)!$ has $(n-1)!$ terms. Also, if C has k elements and D has $n-k$ elements then $(C-a_1)!(D-a_1)!$ has $(k-1)!(n-k-1)!$ terms. Further, there are $\binom{n-2}{k-1}$ pairs C,D with k and $n-k$ elements respectively, since $a_1\in C$ and $a_2\in D$. Hence the number of terms on the right hand side is

$$\begin{aligned}\sum_{k=1}^{n-1}\binom{n-2}{k-1}(k-1)!(n-k-1)! &= \sum_{k=1}^{n-1}(n-2)!\\ &= (n-1)!\end{aligned}$$

as desired.

In looking at the individual terms recall that

$$(B-a_1)!=1(1+q^{a_2-a_1})\cdots(1+q^{a_2-a_1}+\cdots+q^{a_{n-1}-a_1})$$

and that each term in the expansion of $(B-a_1)!$ is obtained by taking the product of one summand from each factor. Furthermore, the factors correspond to the non-minimal elements of B. We must first determine the C and D corresponding to this term. By definition, $a_1\in C$ and $a_2\in D$. The summand from the factor corresponding to a_3 can be either 1 or $q^{a_2-a_1}$. In the former case put $a_3\in C$ and in the latter put $a_3\in D$. Continue putting elements into C or D in increasing order using the following criterion: If the summand for a_i is $q^{a_j-a_1}$ (so $i>j$), put a_i in the same block as a_j. Clearly this uniquely determines a partition $B=C\uplus D$. Also, by definition, the factor corresponding to a_i in $(C - a_1)!$ (or $(D - a_1)!$) contains a unique summand of the form $q^{a_j - a_1}$. Hence the product of these summands is the term in question. ■

We are almost ready to give the q-analog of the the Möbius function of Π_n. First, however, we must define a related function ϕ on Π_n. Later we will show that if x is an element in one of our q-partition posets corresponding to a partition $\pi \in S(\hat{n}, k)$ via a type map, then

$$\mu(x) = (-1)^{n-k}\phi(\pi). \tag{17}$$

If $\pi = B_1/B_2/\ldots/B_k$ is a partition of n then define

$$\phi(\pi) = \prod_{i=1}^{k} q^{\operatorname{int} B_i}(B_i - m_i)! \tag{18}$$

It is clear that substituting $q = 1$ into ϕ yields $|\mu(\pi)|$ as given in Theorem 1.2. Also, if $B_i = \{a_1 < a_2 < \ldots < a_l\}$ then $q^{\operatorname{int} B_i}(B_i - m_i)!$ can be considered as a product with a factor corresponding to each element of $B_i \setminus \{a_1\}$. The *factor corresponding to* a_j, $j \geq 2$, is

$$q^{\operatorname{int} a_j}(1 + q^{a_2 - a_1} + \cdots + q^{a_{j-1} - a_1}).$$

In fact, we can extend this definition to a_1 by letting it correspond to a factor of 1. Hence $\phi(\pi)$ can be thought of as a product with a factor corresponding to each element of $\hat{n}$. This will be used to simplify the proof of the next theorem, a result which will be critical in proving that the Möbius values in our q-analogs cancel properly. In it, a *singleton block* of a partition is one containing only one element.

Theorem 4.3 *Let $\pi = B_1/B_2/\ldots/B_k$ be a partition and let B be the first non-singleton block of π with $B = \{a_1 < \ldots < a_l\}$, then*

$$\phi(\pi) = \sum_{\sigma} q^{\operatorname{nin} \sigma \setminus \pi}\phi(\sigma) \tag{19}$$

where the sum is taken over all refinements, σ, of π which split B into two parts, $B = C \uplus D$ with $a_1 \in C$ and $a_2 \in D$, and keep all other parts the same.

Proof. First consider any element b not in block B. The factor in $\phi(\pi)$ corresponding to b involves $q^{\operatorname{int}_\pi b}$ but the rest of the factor depends only on which elements are in the same block with b. Since this part of the factor does not change in the σ's under consideration, we may cancel it from both sides of equation (19). In order to cancel $q^{\operatorname{int}_\pi b}$ we need

$$q^{\operatorname{int}_\pi b} = q^{\operatorname{int}_\sigma b} q^{\operatorname{nin}_{\sigma \setminus \pi} b}$$

for these σ. But this follows directly from Lemma 4.1.

Hence, after cancelling these factors we are left with the following to be verified:

$$q^{\operatorname{int}_\pi B}(B-a_1)! = \sum q^{\operatorname{nin}_{\sigma\setminus\pi} C} q^{\operatorname{int}_\sigma C}(C-a_1)!\ q^{\operatorname{nin}_{\sigma\setminus\pi} D} q^{\operatorname{int}_\sigma D}(D-a_2)!$$

where the sum is taken over all partitions of the set $\{a_1, a_2, \ldots, a_l\}$ into two parts, C and D, where $a_1 \in C$ and $a_2 \in D$.

First we can simplify this expression by noting that $\operatorname{int}_\pi B = 0$ since B is the first non-singleton block in π. Also, since C will now be the first non-singleton block in σ we see that $\operatorname{nin}_{\sigma\setminus\pi} C = 0$ and $\operatorname{int}_\sigma C = 0$. Further, any element $b \in D, b \neq a_2$ will cause either a new non-inversion or interference with every element, a such that $a_1 < a < a_2$ since b is now in a later block than a. Hence $\operatorname{nin}_{\sigma\setminus\pi} D + \operatorname{int}_\sigma D = (|D|-1)(a_2-a_1)$. Thus we must verify that

$$(B-a_1)! = \sum (C-a_1)!(D-a_2)!\ q^{(|D|-1)(a_2-a_1)} \tag{20}$$

where the sum is taken over C and D as described above.

For the proof of (20) let $D = \{a_2, a_{i_2}, a_{i_3}, \ldots, a_{i_m}\}$; then

$$(D-a_2)! = 1(1+q^{a_{i_2}-a_2})\cdots(1+q^{a_{i_2}-a_2}+\cdots+q^{a_{i_{m-1}}-a_2}).$$

Thus there are $|D|-1$ factors (counting the 1). Further, there are $(|D|-1)$ copies of $q^{a_2-a_1}$ so we can multiply each factor by $q^{a_2-a_1}$ to change all (a_i-a_2)'s to (a_i-a_1)'s. Thus

$$(D-a_2)!\ q^{(|D|-1)(a_2-a_1)} = (D-a_1)!$$

So, the original equation reduces to

$$(B-a_1)! = \sum (C-a_1)!(D-a_1)!$$

which is true by Lemma 4.2. ■

We need one last lemma for the verification that the q-Möbius function will be well behaved.

Lemma 4.4 *Let $\pi = B_1/B_2/\ldots/B_k$ be a partition and let $B = \{a_1, a_2, \ldots, a_l\}$ be the first non-singleton block of π. Then let $\sigma = C_1/C_2/\ldots/C_m$ be any refinement of π such that a_1 and a_2 are still in the same block, say $\{a_1, a_2\} \subseteq C$. Further, let $\tau = D_1/D_2/\ldots/D_{m+1}$ be a refinement of σ such that $C = D \cup D'$ with $a_1 \in D$ and $a_2 \in D'$ and other blocks remain the same. Then*

$$\operatorname{nin} \tau \setminus \pi = \operatorname{nin} \tau \setminus \sigma + \operatorname{nin} \sigma \setminus \pi.$$

Proof. Since B is the first non-singleton block of π, it is easy to see that all the blocks preceding C in σ, or D in τ are singletons. Suppose first that $(a,b) \in \operatorname{Nin}\tau \setminus \pi$. This gives the following possibilities.

1. $\underline{a = \min D'}$. Then $a \in C \setminus \min C$ so $(a,b) \in \operatorname{Nin}\tau \setminus \sigma$.

2. $\underline{b \in D'}$. Then a must be the minimum of a block weakly to the right of D and strictly to the left of D', since $(a,b) \notin \operatorname{Nin}\pi$. Since $b \in C$, we have $(a,b) \in \operatorname{Nin}\tau \setminus \sigma$.

3. $\underline{a, b \notin D'}$. Then the blocks containing a and b are in the same relative positions in τ and σ. Also, $(a,b) \notin \operatorname{Nin}\pi$. Thus $(a,b) \in \operatorname{Nin}\sigma \setminus \pi$.

Hence we have shown that

$$\operatorname{Nin}\tau \setminus \pi \subseteq \operatorname{Nin}\tau \setminus \sigma \uplus \operatorname{Nin}\sigma \setminus \pi.$$

For the reverse inclusion, we have two cases.

(i) $\underline{(a,b) \in \operatorname{Nin}\tau \setminus \sigma}$. Then either a or b must be in D'. (If not, then $(a,b) \in \operatorname{Nin}\sigma$.) If $a = \min D'$ then it is no longer a minimum in π. Thus $(a,b) \in \operatorname{Nin}\tau \setminus \pi$. If $b \in D'$ is not a minimum, then a must be in a block weakly right of D and strictly left of D' (since $(a,b) \notin \operatorname{Nin}\sigma$). So again $(a,b) \notin \operatorname{Nin}\pi$ and $(a,b) \in \operatorname{Nin}\tau \setminus \pi$.

(ii) $\underline{(a,b) \in \operatorname{Nin}\sigma \setminus \pi}$. First of all, $a \notin D'$. (If $a \in D'$ then a is not a minimum in σ, contradicting $(a,b) \in \operatorname{Nin}\sigma$.) Also $b \notin D'$. (If not, then $b \in C$ in σ and $b \in B$ in π. So $(a,b) \in \operatorname{Nin}\sigma$ implies $(a,b) \in \operatorname{Nin}\pi$, a contradiction.) Thus the blocks of a and b are in the same relative position in σ and τ. Hence $(a,b) \in \operatorname{Nin}\tau \setminus \pi$. ■

5 The family $\mathcal{P}_n(q)$

The time has come to describe our q-partition posets. A poset, P, is of *rank* n if it is ranked, and all maximal chains of greatest length have length n. A *type map* is an order-preserving map $T : P \to \Pi_n$. We say that $x \in P$ has *type* π if $T(x) = \pi$. Now define $\mathcal{P}_n(q)$ to be the family of all posets, P, of rank $n-1$ such that there is a type map $T : P \to \Pi_n$ satisfying the following two conditions.

1. The number of elements of P of type π is $q^{\operatorname{nin}\pi}$ for all $\pi \in \Pi_n$.

2. Given $x \in P$ of type π, then the number of $y \leq x$ of type σ is $q^{\operatorname{nin}\sigma\setminus\pi}$.

Conditions 1 and 2 will ensure that the Whitney numbers $W_k(P)$ and $w_k(P)$, respectively, will turn out to be q-Stirling numbers.

It follows immediately from the results in Sections 2 and 3 that we have already seen two posets in $\mathcal{P}_n(q)$.

Theorem 5.1 *The posets $\Pi_n(q)$ and $P_n(q)$ are elements of $\mathcal{P}_n(q)$.* ■

We now show that the the Möbius function for any $P \in \mathcal{P}_n(q)$ is given as in equation (17).

Theorem 5.2 *Suppose $P \in \mathcal{P}_n(q)$. If $x \in P$ has type $\pi = B_1/\ldots/B_k$ then*

$$\mu(x) = (-1)^{n-k}\phi(\pi) \tag{21}$$

Proof. If $\pi = 1/2/\ldots/n$ then $\operatorname{nin}\pi = 0$. So there is only $q^0 = 1$ element x of type π, by condition 1 in the definition of $\mathcal{P}_n(q)$. Thus x must be a unique minimum of P and hence $\mu(x) = 1$. But the definition of ϕ shows that the right side of equation (21) is also 1, giving agreement in this case.

Now let $\pi = B_1/B_2/\ldots/B_k$ be an element of Π_n. Define a function on P by $f(x) = (-1)^{n-k}\phi(\pi)$. We wish to show that

$$\sum_{y\le x} f(y) = 0$$

for any $x \neq \hat{0}$. Then f will satisfy the recurrence for the Möbius function and by uniqueness we will have $f = \mu$.

Now use condition 2 of the definition of $\mathcal{P}_n(q)$. For any refinement σ of π there are $q^{\operatorname{nin}\sigma\backslash\pi}$ elements $y \in P$ which are below x and of type σ. If σ has l blocks, then

$$\begin{aligned}\sum_{y\le x} f(y) &= \sum_{\substack{y\le x\\ \sigma=T(y)}} (-1)^{n-l}\phi(\sigma)\\ &= \sum_{\sigma\le\pi} (-1)^{n-l} q^{\operatorname{nin}\sigma\backslash\pi}\phi(\sigma). \end{aligned} \tag{22}$$

Let B and C be defined as in Lemma 4.4. Then for each σ which keeps a_1 and a_2 in the same block Theorem 4.3 says (with a change of notation) that

$$\phi(\sigma) = \sum_{\tau} q^{\operatorname{nin}\tau\backslash\sigma}\phi(\tau)$$

where the sum is taken over all refinements τ of σ which split C into two parts with a_1 and a_2 separated but keep all other parts the same. Hence by Lemma 4.4

we have

$$\begin{aligned} q^{\operatorname{nin}\sigma\backslash\pi}\phi(\sigma) &= \sum_{\tau} q^{\operatorname{nin}\tau\backslash\sigma} q^{\operatorname{nin}\sigma\backslash\pi}\phi(\tau) \\ &= \sum_{\tau} q^{\operatorname{nin}\tau\backslash\pi}\phi(\tau) \end{aligned}$$

for refinements τ of σ which split a_1 and a_2 but keep other blocks the same. Thus every term in (22) cancels out since every refinement of π (including π itself) either has a_1 and a_2 in the same block and so cancels with even finer refinements or has a_1 and a_2 separated and so cancels as part of some coarser refinement. Hence we have

$$\sum_{y \leq x} f(y) = 0$$

as desired so $\mu(x) = (-1)^{n-k}\phi(\pi)$. ■

We end this section with a q-analog of Theorem 1.3.

Theorem 5.3 *The Whitney numbers for any $P \in \mathcal{P}_n(q)$ are*

$$w_k(P) = s[n, n-k] \quad \textit{and} \quad W_k(P) = S[n, n-k].$$

Proof. We will first prove the statement about the Whitney numbers of the second kind. Because the type map is order-preserving, the elements at rank $n-k$ in $P \in \mathcal{P}_n(q)$ are exactly those which map to partitions π of rank $n-k$ in Π_n. Each such π has k blocks and $q^{\operatorname{nin}\pi}$ elements in its preimage. Thus

$$W_{n-k}(P) = \sum_{\pi = B_1/\ldots/B_k} q^{\operatorname{nin}\pi} = S[n,k]$$

by Theorem 1.1.

Now we will verify the values of the Whitney numbers of the first kind.

$$\begin{aligned} w_{n-k}(P) &= \sum_{\operatorname{rk} x = n-k} \mu(x) \\ &= \sum_{\pi = B_1/\ldots/B_k} (-1)^{n-k} q^{\operatorname{nin}\pi}\phi(\pi) \\ &= (-1)^{n-k}\left[\sum_{B_k = \{n\}} q^{\operatorname{nin}\pi}\phi(\pi) + \sum_{B_k \neq \{n\}} q^{\operatorname{nin}\pi}\phi(\pi)\right] \\ &= (-1)^{n-k}\left[\sum_{\pi' \in S(\widehat{n-1}, k-1)} q^{\operatorname{nin}\pi'}\phi(\pi') + \sum_{\substack{\pi \text{ with} \\ \pi' \in S(\widehat{n-1}, k)}} r_{\pi/\pi'} q^{\operatorname{nin}\pi'}\phi(\pi')\right] \end{aligned}$$

where π' is π with n deleted and

$$r_{\pi/\pi'} = q^{\mathrm{nin}\,\pi}\phi(\pi)/q^{\mathrm{nin}\,\pi'}\phi(\pi').$$

Suppose the block containing n in π is $B = \{a_1 < a_2 < \ldots < a_m < n\}$. Then

$$\begin{aligned} r_{\pi/\pi'} &= q^{\mathrm{nin}\,n}q^{\mathrm{int}\,n}(1 + q^{a_2-a_1} + \cdots + q^{a_m-a_1}) \\ &= q^{a_1-1}(1 + q^{a_2-a_1} + \cdots + q^{a_m-a_1}) \\ &= q^{a_1-1} + q^{a_2-1} + \cdots + q^{a_m-1} \end{aligned}$$

Holding π' fixed and summing over all π obtained by adding n to a block of π' yields

$$\sum_{\pi} r_{\pi/\pi'} = \sum_{C\in\pi'}\left(\sum_{a\in C} q^{a-1}\right) = \sum_{a=1}^{n-1} q^{a-1} = [n-1].$$

Note that this sum does not depend on π'. So, by induction and the definition of $s[n,k]$, we have

$$\begin{aligned} w_{n-k}(P) &= (-1)^{n-k}\left[\sum_{\pi'\in S(\widehat{n-1},k-1)} \phi(\pi') \;+\; [n-1] \sum_{\pi'\in S(\widehat{n-1},k)} \phi(\pi')\right] \\ &= s[n-1,k-1] - [n-1]s[n-1,k] \\ &= s[n,k] \quad \blacksquare \end{aligned}$$

6 Comments and open questions

Here are some remarks about the foregoing material. It will be convenient to use the notation $\pi = /B_1/B_2/\ldots/B_k/$ for the partition of $\hat{n} \supseteq \uplus_i B_i$ such that all elements of $\hat{n} \setminus \uplus_i B_i$ are in singleton blocks.

(1) Even though $\Pi_n(q)$ and $P_n(q)$ were constructed differently, they might still be isomorphic as posets. The next result rules out this possibility.

Proposition 6.1 *For $n \geq 4$ and $q \geq 2$ we have $\Pi_n(q) \not\cong P_n(q)$.*

Proof. Suppose, to the contrary that there is an isomorphism $f : \Pi_n(q) \to P_n(q)$. Then f is also a graph-theoretic isomorphism of the Hasse diagrams of these posets. In particular, the bipartite graphs induced by the elements at ranks 1 and 2 in these Hasse diagrams are also isomorphic. Furthermore, f must also preserve the Möbius function values of these elements.

We claim that f carries the elements of type $/n-1\ n/$ in $\Pi_n(q)$ to those of the same type in $P_n(q)$. An $\ell \in \Pi_n(q)$ of type $/n-1\ n/$ is adjacent to an element of type $\pi = /1\ n-1\ n/$. But π is the unique type at rank 2 having Möbius value $1+q^{n-2}$. So $f(\ell)$ must be an element of rank 1 of type σ where $\sigma < \pi$. Thus the possible choices for σ are $/1\ n-1/$ or $/1\ n/$ or $/n-1\ n/$. Now, for $n \geq 4$, elements of type $/n-1\ n/$ are also adjacent to those of type $/1\ 2/n-1\ n/$, which have μ-function q. However, elements of type $/1\ n-1/$ or $/1\ n/$ have no neighbors with μ-function q. Hence, f preserves the type of the $/n-1\ n/$ elements.

Now consider the elements of type $/n-2\ n-1\ n/$. These are adjacent to those of type $/n-1\ n/$. Furthermore, they are the only neighbors of the $/n-1\ n/$ elements having Möbius value $1+q$. This fact, together with what we proved in the previous paragraph, shows that f must restrict to an isomorphism of the bipartite subgraphs induced by the elements of types $/n-1\ n/$ and $/n-2\ n-1\ n/$ in $\Pi_n(q)$ and $P_n(q)$.

We will now show that this is an impossible situation. In $\Pi_n(q)$ this subgraph consists of q^{n-3} copies of the complete bipartite graph $K_{q,q^{n-3}}$. To see this, note that the vertices of type $/n-1\ n/$ in a given component consist of all ℓ with last line $(a_1, \ldots, a_{n-2}, 1, 0)$, where $a_2, \ldots, a_{n-2}$ are fixed and a_1 varies over F_q. The corresponding vertices of type $/n-2\ n-1\ n/$ have last two lines

$$\begin{array}{c}(b_1, \ldots, b_{n-3}, 1, 0, 0)\\(a_2, \ldots, a_{n-2}, 1, 0, 0)\end{array}$$

where $b_1, \ldots b_{n-3}$ vary over F_q. This means that, in the subgraph of $\Pi_n(q)$, the neighborhoods of any two vertices of type $/n-1\ n/$ are either equal or disjoint.

However, we can construct two vertices of type $/n-1\ n/$ in $P_n(q)$ whose neighborhoods are neither equal nor disjoint. Let

$$v_i = (\infty, \ldots, \infty, 1, 0, \ldots, 0)$$

where the pivot 1 is in the i^{th} position. (Remember that the positions are numbered $2, \ldots, n$.) Also let

$$w = (\infty, \ldots, \infty, 1, 1)$$

and consider the matrices

$$M = \begin{pmatrix} v_2 \\ \vdots \\ v_{n-2} \\ v_{n-1} \end{pmatrix} \qquad \text{and} \qquad N = \begin{pmatrix} v_2 \\ \vdots \\ v_{n-2} \\ w \end{pmatrix}.$$

Then both M and N are adjacent to

$$\begin{pmatrix} v_2 \\ \vdots \\ v_{n-2} \end{pmatrix}.$$

But M is adjacent to

$$\begin{pmatrix} v_2 + 1 \cdot v_{n-1} \\ \vdots \\ v_{n-2} \end{pmatrix}.$$

while N is not. This contradiction ends the proof of the proposition. ■

(2) Although Π_n is a lattice, the elements of $\mathcal{P}_n(q)$ are not as the following result shows.

Proposition 6.2 *If $P \in \mathcal{P}_n(q)$ for $n \geq 4$ and $q \geq 2$ then P is not a lattice.*

Proof. Since $n \geq 4$, we can consider the partitions

$$\pi = /1/2/34/, \qquad \sigma = /134/2/, \qquad \tau/12/34/.$$

By condition 1 in the definition of $\mathcal{P}_n(q)$,

$$|T^{-1}(\pi)| = q^2, \qquad |T^{-1}(\sigma)| = 1, \qquad |T^{-1}(\tau)| = q.$$

Also, $\operatorname{nin} \pi \setminus \sigma = 2$. So, by condition 2, every element of type π lies below the unique element of type σ. Finally, $\operatorname{nin} \pi \setminus \tau = 1$, so q^1 elements of type π lie below each element of type τ. Since $q \geq 2$, we can choose $x, y \in T^{-1}(\pi)$ both of which are covered by the element of type σ and one element of type τ. Hence x and y cannot have a join. ■

(3) Recently Simion [Sim pr] has developed a general theory of q-analogs of posets that includes our constructions as a special case. Using her machinery, one can show our posets are EL-shellable [BGS 82]. Thus one can get a combinatorial explanation for the factors in (18) by counting decreasingly labeled chains.

References

[BGS 82] A. Björner, A. M. Garsia and R. P. Stanley, An introduction to Cohen-Macaulay partially ordered sets, in "Ordered Sets," I. Rival, ed., Reidel, Dordrecht, 1982, 583–615.

[Car 33] L. Carlitz, On Abelian fields, *Trans. Amer. Math. Soc.* **35** (1933), 122–136.

[Car 48] L. Carlitz, q-Bernoulli numbers and polynomials, *Duke Math. J.* **15** (1948), 987–1000.

[Dem 92] K. J. Dempsey, "q-Analogs and Vector Spaces," Ph. D. thesis, Michigan State University, East Lansing, 1992.

[Dow 73] T. A. Dowling, A q-analog of the partition lattice, in "A Survey of Combinatorial Theory," J. N. Srivastava et al., eds., North-Holland Pub. Co., Amsterdam, 1973, 101–115.

[Ges 82] I. M. Gessel, A q-analog of the exponential formula, *Discrete Math.* **40** (1982), 69–80.

[Gou 61] H. W. Gould, The q-Stirling numbers of the first and second kinds, *Duke Math. J.* **28** (1961), 281–289.

[Mil 78] S. C. Milne. A q-analog of restricted growth functions, Dobinski's equality, and Charlier polynomials, *Trans. Amer. Math. Soc.* **245** (1978), 89–118.

[Mil 82] S. C. Milne, Restricted growth functions, rank row matchings of partition lattices, and q-Stirling numbers, *Adv. in Math.* **43** (1982), 173–196.

[Sim pr] R. Simion, On q-analogues of partially ordered sets, preprint.

[Sag 91] B. E. Sagan, A maj statistic for set partitions, *European J. Combin.* **12** (1991), 69–79.

[Sta 86] R. P. Stanley, "Enumerative Combinatorics, Vol. 1," Wadsworth & Brooks/Cole, Pacific Grove, CA, 1986.

[W-W 91] M. Wachs and D. White, p, q-Stirling numbers and set partition statistics, *J. Combin. Theory Ser. A* **56** 1991, 27–46.