Deformation of Chains via a Local Symmetric Group Action

Patricia Hersh*
Department of Mathematics, Room 2-588
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
hersh@math.mit.edu

Submitted: May 13, 1998; Accepted: March 3, 1999.

AMS Subject Classification: 05E25, 06A07.

Abstract

A symmetric group action on the maximal chains in a finite, ranked poset is local if the adjacent transpositions act in such a way that $(i, i+1)$ sends each maximal chain either to itself or to one differing only at rank i. We prove that when S_n acts locally on a lattice, each orbit considered as a subposet is a product of chains. We also show that all posets with local actions induced by labellings known as R^*S-labellings have symmetric chain decompositions and provide R^*S-labellings for the type B and D noncrossing partition lattices, answering a question of Stanley.

1 Introduction

A symmetric group action on the maximal chains in a finite, ranked poset was defined by Stanley in [St2] to be local if for each i, the adjacent transposition $s_i = (i, i+1)$ sends each maximal chain either to itself or to one differing from it only at rank i.

*This work was supported by a Hertz Foundation Graduate Fellowship.

There is a correspondence between rhombic tilings of a planar region and equivalence classes of reduced expressions for a permutation up to commutation. This naturally translates symmetric group structure to poset structure when S_n acts locally on the maximal chains in a poset. We begin by reviewing this correspondence which is thoroughly examined in [El] because we it will allow us to explain why orbits of local symmetric group actions on lattices are always products of chains.

When a permutation w is written as a product of adjacent transpositions $w = s_{a_1}s_{a_2}\dots s_{a_l}$ with l as small as possible, such a product is called a reduced expression for w. To obtain a rhombic tiling from this, begin with a vertical path consisting of $n+1$ nodes; as one reads off each successive adjacent transposition s_{a_i} in a reduced expression, draw a new node to the right of the current node of rank a_i, and attach this new node to the nodes of rank $a_i \pm 1$ in the current path to obtain a new path. The resulting region is bounded on the left by the initial path, on the right by the final path, and is tiled by quadrilaterals. These quadrilaterals may be replaced by rhombi by appropriately adjusting line segment slopes. Two reduced expressions differing only by commutation relations give rise to the same rhombic tiling. Applying a braid relation $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$ to a reduced expression amounts to a substitution within a tiling as in Figure 1. Any two

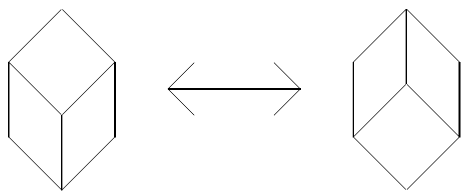

Figure 1: The relation $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$ in terms of tilings

reduced expressions for the same permutation give rise to rhombic tilings which fit in exactly the same planar region. One may obtain any rhombic tiling for a particular region from any other by applying braid relations.

We will use rhombic tilings to record how a maximal chain is deformed under a local symmetric group action by successively applying the adjacent transpositions in a reduced expression for a permutation. If $p_2 = wp_1$, then each reduced expression for w gives rise to a (potentially distinct) way of deforming the maximal chain p_1 to p_2 within a poset. The structure of a poset with a local symmetric group action must allow for all possible ways of deforming one maximal chain to another.

Each rhombic tiling may be viewed as the projection of a discrete 2-dimensional surface S within a hypercube or multi-dimensional box onto a generic plane. Such a surface S may be deformed via braid relations (as in Figure 1) to surfaces coming from other reduced expressions for the same permutation; relations of the form $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$ will take surfaces which include the front three faces of a cube to surfaces which instead includes the back three faces. The surfaces given by the same permutation will have the same boundary. The collection of rhombic tilings for a particular region gives rise to all the minimal discrete surfaces within a multi-dimensional box which have some fixed boundary. This point of view leads us to prove in Section 2 that the maximal chains in an orbit of a local symmetric group action must be arranged in such a way that they form the skeleton of such a multi-dimensional box. Otherwise, braid relations would be violated or an orbit would be incomplete (or both).

This does not, however, imply that each orbit is a product of chains since the nodes in such a skeleton need not all be distinct. In Section 2, we prove that the nodes are distinct when the poset is a lattice and conclude that the orbits in lattices are products of chains. Sections 3 and 4 examine local actions induced by labellings known as R^*S-labellings: in the former section we prove that all posets with R^*S-labellings have symmetric chain decompositions, while the latter provides R^*S-labellings for the interpolating BD noncrossing partition lattices.

2 Orbit characterization

Simion and Stanley have shown in [SS] that the Frobenius characteristic of a local symmetric group action on an orbit is always a complete symmetric function. Theorem 3 will provide a more geometric proof of this result in order to show how orbits are realized within posets. This will allow us to characterize the orbits of local symmetric group actions on lattices in Theorem 4, answering a question of Stanley.

Figure 2 gives an example of how the situation differs between posets and lattices. Identifying the nodes labelled $(0,3)$ and $(3,0)$ within a product of two 4-chains, yields a poset with a local symmetric group action with three orbits. One orbit consists of the maximal chains from the original product of chains before identification. Two new maximal chains are introduced, one of which is depicted by the shaded lines in Figure 2. These maximal chains due to crossover give rise to two trivial orbits.

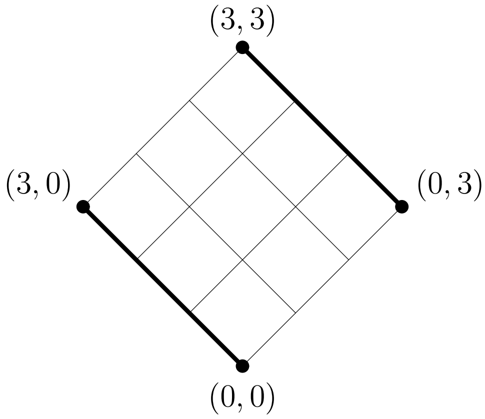

Figure 2: A product of chains with node identification

This example depends on the fact that the above poset is not a lattice.

Definition 1. *When the symmetric group acts locally on the maximal chains in a poset, then the elements of an orbit subposet are the nodes in any maximal chain within the orbit specifying it. The covering relations are induced by covering relations from the maximal chains in the orbit.*

In lattices, the maximal chains in an orbit subposet turn out to be exactly the maximal chains belonging to the orbit specifying it.

Lemma 2 justifies geometric claims within the proof of Theorem 3.

Lemma 2. *If $s_i(p) \neq p$ and $s_{i+1}(p) = p$, then $s_{i+1}(s_i(p)) \neq s_i(p)$. Similarly, if $s_{i+1}(p) \neq p$ and $s_i(p) = p$, then $s_i(s_{i+1}(p) \neq s_{i+1}(p)$.*

PROOF. If $s_{i+1}(s_i(p)) = s_i(p)$ and $s_{i+1}(p) = p$, then

$$\begin{aligned} s_i(p) &= s_{i+1}(s_i(p)) \\ &= s_{i+1}s_i(s_{i+1}(p)) \\ &= s_i(s_{i+1}(s_i(p))) \\ &= s_i(s_i(p)) \\ &= p. \end{aligned}$$

The second assertion follows similarly. □

In Theorem 3 we will define a map ϕ from maximal chains in a poset to lattice paths in $\mathbb{Z}^n$. Lemma 2 implies that whenever $im(\phi)$ includes two lattice paths involving segments **abdf** and **acdf**, (in Figure 3) respectively, and otherwise agreeing, then $im(\phi)$ will also include a lattice path through **acef** which otherwise agrees with both these paths. No assumption is made about whether **bd** is perpendicular or parallel to **df**.

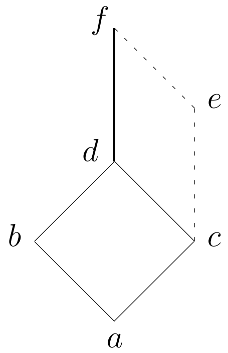

Figure 3: Building an orbit

Theorem 3. *If S_n acts locally on the maximal chains in a poset, then the Frobenius characteristic of the action is an h-positive symmetric function.*

PROOF. We will prove that the local symmetric group action on any orbit is isomorphic to a local action on some product of chains $C_{\lambda_1+1} \times \cdots \times C_{\lambda_k+1}$, since this action will have Frobenius characteristic h_λ.

Let us refer to maximal chains in a poset P as P-chains. We claim that any orbit may be embedded by a map ϕ into the lattice $\mathbb{Z}^n$ in such a way that poset rank is encoded as sum of coordinates and P-chains are sent to lattice paths within $\mathbb{N}^n$. We will define ϕ in such a way that $im(\phi)$ will be the collection of all minimal lattice paths from the origin to a particular endpoint in $\mathbb{N}^n$. Furthermore, s_i will act nontrivially on a P-chain p whenever the segment of the lattice path $\phi(p)$ from rank $i-1$ to rank i is perpendicular to the segment from rank i to $i+1$. When path segments are labelled by lattice basis vectors, then applying an adjacent transposition will amount to swapping a lattice path with one which has the two consecutive labels swapped.

We define ϕ by choosing a P-chain p and specifying how to embed wp into $\mathbb{N}^n$ for each $w \in S_n$. The embedding will be based on a choice of reduced expression for w, but we will check that all reduced expressions for the same permutation w yield the same lattice path $\phi(wp)$. To conclude that ϕ is well-defined, we will also need to show that $\phi(w_1p) = \phi(w_2p)$ whenever $w_1p = w_2p$, using the definition of local action.

If $s_i(p) = p$ for all $i < a_1$ and $s_{a_1}(p) \neq p$, then let the lattice path $\phi(p)$ begin with a segment from $(0, \ldots, 0)$ to $(a_1, 0, \ldots, 0)$. The lattice path $\phi(p)$

will change direction at rank i for each i such that $s_i(p) \neq p$. In particular, if s_{a_2} acts nontrivially on p and all intermediate s_i act trivially on p, then $\phi(p)$ includes the segment from $(a_1, 0, \ldots, 0)$ to $(a_1, a_2 - a_1, 0, \ldots, 0)$. At this point, we may specify how $\phi(wp)$ is related to $\phi(p)$ for any w which only involves the adjacent transpositions $s_1, \ldots, s_{a_2-1}$. If $s_j p_1 = p_2 \neq p_1$ for a P-chain p_1 which has already been embedded up to rank $j+1$, then p_2 is embedded up to rank $j+1$ by replacing the node of rank j in $\phi(p_1)$ with the only other node of rank j in $\mathbb{N}^n$ which together with the rest of $\phi(p_1)$ gives a lattice path. In this way, the embedding of p up to rank a_2 locally gives rise to every possible discrete path of minimal length from the origin to $(a_1, a_2 - a_1, 0, \ldots, 0)$; first one obtains $\phi(s_{a_1}(p))$, and repeated application of Lemma 2 yields all minimal length lattice paths from $(0, \ldots, 0)$ to $(a_1, a_2 - a_1, 0, \ldots, 0)$. These paths may be sequentially embedded in many different orders, but the commutation relations $s_i s_j = s_j s_i$ for $|j - i| \geq 2$ force all choices to be equivalent.

The direction in which to extend $\phi(p)$ to rank $a_2 + 1$ is determined by how s_{a_2} acts upon the P-chains with image under ϕ passing through the lattice point $(a_1, a_2 - a_1, 0, \ldots, 0)$ which also agree with $\phi(p)$ afterwards. The edge out of $(a_1, a_2 - a_1, 0, \ldots, 0)$ in $\phi(p)$ needs to be perpendicular to exactly those segments into $(a_1, a_2 - a_1, 0, \ldots, 0)$ which belong to lattice paths which are acted upon nontrivially by s_{a_2}, and which also include the given segment out of $(a_1, a_2 - a_1, 0, \ldots, 0)$.

Lemma 2 implies that at each step of the embedding of p, the next segment of $\phi(p)$ should be perpendicular to all but at most one of the lattice path edges leading into this new segment, so embedding is feasible. In this fashion we may define $\phi(p)$. Each time $\phi(p)$ changes direction, we repeatedly apply Lemma 2 just as we did at rank a_1 to obtain lattice paths of the form $\phi(wp)$. The relations $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$ imply that when three consecutive segments of some $\phi(wp)$ are all perpendicular, six lattice paths result all belonging to $im(\phi)$, and the restriction of these lattice paths to the interval form the skeleton of a cube.

Repeated application of Lemma 2 and braid relations thus yields every minimal lattice path from the origin to the endpoint of $\phi(p)$ as the image of some P-chain, so ϕ will be onto. We need only show that any pair of distinct lattice paths $\alpha, \beta \in im(\phi)$ come from distinct P-chains to insure that ϕ is well-defined. Let v_1 and v_2 be nodes in $\mathbb{Z}^n$ where α and β first differ. There must also exist lattice paths $\gamma, \gamma' \in im(\phi)$ containing v_1 and v_2, respectively, which otherwise agree with each other. From the definition of ϕ it follows that γ and γ' are the images of distinct P-chains q, q' which satisfy $q' = s_i(q)$ for $i = rank(v_1)$. Hence, v_1 and v_2 must be the images of distinct poset elements of rank i, implying α and β are the images of distinct

P-chains, so ϕ is well-defined. Our definition of ϕ insures that ϕ is injective, since $\phi(p) \neq \phi(wp)$ whenever $p \neq wp$.

The local S_n-action on the orbit is thus an action well-known to have Frobenius characteristic h_λ, as desired. □

The argument we present next was gleaned from a more complicated proof involving the correspondence between rhombic tilings and commutation classes of reduced expressions for a permutation.

Theorem 4. *If S_n acts locally on a lattice, then each orbit is a product of chains.*

PROOF. We begin by proving that a poset obtained from a product of two chains by identifying two of its nodes cannot be a lattice. After this, we will show how to reduce the proof of the theorem to this case. We assume throughout that there is no node identification at rank 1, because we dealt with this possibility while proving ϕ was well-defined in Theorem 3.

Consider a product of two chains, each of which has rank r. Let us identify $a = (r, 0)$ with $b = (0, r)$ and assume there is no node identification below rank r. Suppose this poset is a lattice. We use induction on j to show that $(j, 1) \leq a$ for all $j < r$. As the base case, observe that $(0, 1) \leq a$ since $a = b = (0, r)$. If $(j, 1) \leq (r, 0)$ for some $j \geq 0$, then $(j, 1) \vee (j + 1, 0) \leq (r, 0)$ for $j + 1 \leq r$. Since $(j + 1, 0) \leq (j + 1, 1)$ and $(j, 1) \leq (j + 1, 1)$ and rank $(j + 1, 1) =$ rank $(j, 1) + 1$, note that $(j, 1) \vee (j + 1, 0) = (j + 1, 1)$ in the poset. The definition of join thus implies $(j + 1, 1) \leq (r, 0) = a$ whenever $(j, 1) \leq (r, 0)$ for $j + 1 \leq r$. By induction, $(r - 2, 1) \leq a$, so $a \geq (r - 2, 1) \vee (r - 1, 0) = (r - 1, 1)$, a contradiction.

There is one somewhat subtle point to be addressed in the way we will show a poset is not a lattice by restricting to some subposet and showing this cannot be a lattice. When we assume a poset is a lattice, we need to be careful about whether the join of two subposet elements also lies in the subposet. In the above induction, we only deal with joins $a \vee b$ of rank one more than the rank of a and b, so this must be the join of a and b in any lattice containing the above as a subposet.

Now consider any product of chains with nodes a and b of rank r identified and with no node identification below rank r. Choose a maximal chain p_1 through the node a and a maximal chain p_2 through b (before identification), and then restrict attention to the nodes in some deformation of p_1 to p_2. We choose p_1 and p_2 so that the number of adjacent transpositions needed to deform p_1 to p_2 is as small as possible. If we let $a = (a_1, \ldots, a_k)$ and

$b = (b_1, \ldots, b_k)$, using the coordinates given by the product of chains structure, then p_1 and p_2 both contain the node $(\min(a_1, b_1), \ldots, \min(a_k, b_k))$ and agree below this node. Furthermore, a minimal deformation will not affect the nodes between $\hat{0}$ and $(\min(a_1, b_1), \ldots, \min(a_k, b_k))$. The nodes above $(\min(a_1, b_1), \ldots, \min(a_k, b_k))$ which occur in a minimal deformation will give rise to a product of two chains, but with a and b identified.

This last observation follows from the fact that the coordinates which increase in travelling from $(\min(a_1, b_1), \ldots, \min(a_k, b_k))$ to a along the maximal chain p_1 are completely disjoint from the set of coordinates which increase in p_2 between $(\min(a_1, b_1), \ldots, \min(a_k, b_k))$ and b. An example is illustrated in Figure 4. The product of two chains comes from interspersing steps in the direction of p_1 with steps in the direction of p_2, while travelling from $(\min(a_1, b_1), \ldots, \min(a_k, b_k))$ to $(\max(a_1, b_1), \ldots, \max(a_k, b_k))$

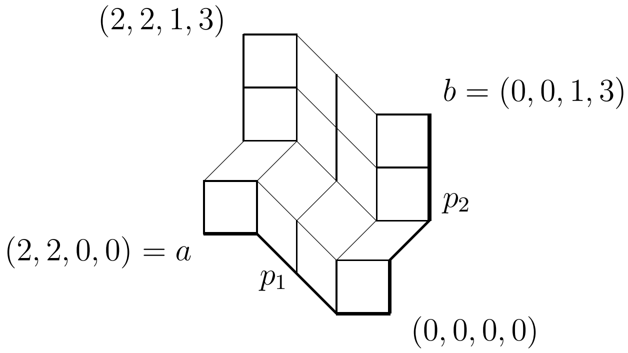

Figure 4: A 2-dimensional surface within a 4-dimensional product of chains

If a and b are identified in any product of chains, they will thus also be identified in a subposet which is a product of two chains, and so the original poset will not be a lattice. □

3 Symmetric boolean decomposition

In this section, we show that posets with R^*S-labellings have symmetric chain decompositions which may be defined in terms of these labellings. The notion of R^*S-labelling was introduced by Simion and Stanley in [SS]. A chain-labelling of a poset is R^* if $\hat{0} \prec u_1 \prec \cdots \prec u_k = u \leq v$ implies there is a unique extension of this chain to a saturated chain from $\hat{0}$ to v with strictly increasing labels between u and v. If a chain-labelling λ induces a local S_n-action on the maximal chains of a poset, and the sequences labelling the maximal chains are all distinct, then λ is an S-labelling; in this case, S_n

acts on sequences of edge labels by permuting the order of the labels, and this induces a local action on the maximal chains with corresponding labels. An S-labelling which is also R^* is an R^*S-labelling.

In the following theorem, we make reference to the unique saturated chain from u to v with increasing labels for any pair $u \leq v$. This is not well-defined for a chain-labelling which is not an edge labelling, but we establish the following convention. When we refer to the unique increasing chain from u to v, we first choose the increasing maximal chain from $\hat{0}$ to u, and then based on this choice we select the resulting increasing saturated chain from u to v.

Theorem 5. *If a finite, ranked poset admits an R^*S-labelling, then the elements may be decomposed into a disjoint union of symmetrically placed boolean lattices.*

PROOF. We define a map ϕ from elements of a finite poset P to symmetrically placed boolean lattices in the poset and show that this map is a decomposition. Let λ be an R^*S-labelling for a poset P of rank n. For each $v \in P$, there are unique saturated chains $\hat{0} = u_0 \prec u_1 \prec \cdots \prec u_k = v$ and $v = v_0 \prec \cdots \prec v_l = \hat{1}$ with strictly increasing labels. Since λ is an S-labelling, there exist u and w such that $\hat{0} \leq u \leq v \leq w \leq \hat{1}$ and rank $w = n-$ rank u with u and w satisfy the following two conditions: first, the set of labels on the unique rising chain from $\hat{0}$ to u is the same as the set of labels on the unique rising chain from w to $\hat{1}$. Second, the set of labels on the rising chain from u to v is disjoint from the set of labels on the rising chain from v to w; each of these sets is also disjoint from the set of labels on the rising chains from $\hat{0}$ to u and from w to $\hat{1}$. This is possible by restricting S_n to acting locally on the saturated chain from $\hat{0}$ to v and likewise on the saturated chain from v to $\hat{1}$ to obtain new saturated chains with all common labels shifted down to below u and up to above w. There is a symmetrically placed boolean lattice $B_{u,w}$ on the interval from u to w. It consists of all nodes reached by restricting S_n to acting locally on the orbit within this interval (u, w) which includes the increasing chain from u to w. Since λ is an R^*-labelling, u and w are uniquely specified, and $v \in B_{u,w}$, so let $\phi(v) = B_{u,w}$.

Note that if $\phi(v_1) = B_{u,w}$ and $v_2 \in B_{u,w}$, then $\phi(v_2) = B_{u,w}$, because the unique increasing chains from $\hat{0}$ to v_2 and from v_2 to $\hat{1}$ may be obtained by taking a maximal chain which includes v_2 in addition to u and w since v_2 belongs to $B_{u,w}$, and then applying a sequence of adjacent transpositions permuting the labels above v_2 and below v_2 separately. Hence, ϕ provides a decomposition. $\square$

Note that R^*S-labellings restrict to intervals, so Theorem 5 also applies to all the intervals in posets with R^*S-labellings.

Corollary 6. *If a finite, ranked poset admits an R^*S-labelling, then it has a symmetric chain decomposition.*

PROOF. Theorem 9 provides a decomposition into symmetrically placed boolean lattices, and each of these has a symmetric chain decomposition. One may find an explicit construction of a symmetric chain decomposition for the boolean lattices in a survery article by Greene and Kleitman [GK], and this article also gives original references (de Bruijn et al., Leeb). □

Reiner gives symmetric chain decompositions for the type B and interpolating BD noncrossing partition lattices in [Re], but his SCD's for interpolating BD noncrossing partition lattices are not an immediate consequence of his recursively defined SCD for type B. In the next section we provide an R^*S-labelling for the type B noncrossing partition lattice, and this is easily shown to restrict to an R^*S-labelling for the interpolating BD noncrossing partition lattices. Theorem 5 leads to Reiner's symmetric chain decomposition for type B, and this restricts to an SCD for the other types since the R^*S-labellings restrict to other types. One may similarly show that other subposet of posets with R^*S-labellings have symmetric chain decompositions.

4 Noncrossing partitions of types B and D

Reiner defined and studied noncrossing partition lattices for the classical reflection groups in [Re]. In this section, we define an R^*S-labelling for the type B, D and interpolating BD noncrossing partition lattices, answering a question raised by Stanley in [St3]. The labelling for type B restricts to one for the other types. This R^*S-labelling is also closely related to the labelling by parking functions for the traditional (type A) noncrossing partition lattice given in [St3].

The type B noncrossing partition is a partition of $\pm 1, \pm 2, \ldots, \pm n$ satisfying the following two conditions. If one places $1, 2, \ldots, n, -1, -2, \ldots, -n$ sequentially about a circle spacing the numbers evenly, and one draws straight lines through the circle between any two numbers belonging to the same component of a partition, then the interior of the circle should have 180 degree rotational symmetry. Furthermore, deleting any two edges that cross should leave the partition unchanged. By convention, the numbers are placed clockwise about the circle. For each component C_i there will be a component $-C_i$

such that $j \in C_i$ if and only if $-j \in -C_i$. In particular, one of the components may satisfy $C = -C$. When such a component exists, this is called the 0-component, denoted by C_0. The noncrossing property precludes the existence of more than one 0-component. The type B noncrossing partition lattice is the poset of type B noncrossing partitions ordered by refinement. The type D noncrossing partition lattice is the restriction to partitions with $C_0 \neq \{\pm i\}$ for $1 \leq i \leq n$. The interpolating BD noncrossing partition lattices arise by choosing a subset $S \subseteq \{1, \ldots, n\}$ and forbidding partitions with $C_0 = \{\pm i\}$ for $i \in S$.

We define an edge-labelling λ for the type B noncrossing partition lattice in terms of covering relations. If v is obtained from u by merging C with $-C$ to form C_0, then C is entirely contained in some semicircle S. In this case, we let $\lambda(u, v) = i$ where $i \in \{1, \ldots, n\}$ and $\pm i$ is the last element of C encountered while travelling clockwise about S. If v is obtained from u by merging two components C_1 and C_2 (and simultaneously merging $-C_1$ with $-C_2$), then one of the following conditions is true (for some choice of which component is C_1 and which is C_2). Either there is some pair of elements both in C_2 such that all elements of C_1 lie between these two elements of C_2, or there is some semicircle S which contains both C_1 and C_2 and such that all of C_2 comes before all of C_1 travelling clockwise about S. In either case, there is some $i \in \{1, \ldots, n\}$ such that $\pm i$ is the last element of C_2 encountered before the first clockwise element of C_1, and then we let $\lambda(u, v) = i$.

For example, listing only blocks with more than one element, the maximal chain

$$\emptyset \prec \pm\{5,6\} \prec \pm\{5,6\}, \{\pm 4\} \prec \pm\{1,3\}, \pm\{5,6\}, \{\pm 4\} \prec \pm\{1,2,3\}, \pm\{5,6\},$$
$$\{\pm 4\} \prec \{\pm 1, \pm 2, \pm 3, \pm 4\}, \pm\{5,6\} \prec \{\pm 1, \pm 2, \pm 3, \pm 4, \pm 5, \pm 6\}$$

is labelled with the parking function $\lambda = (5, 4, 1, 1, 4, 4)$. Figure 5 depicts how components are sequentially merged in this maximal chain; arc labels indicate the order in which components are merged.

We follow Stanley [St3] in referring to sequences in $\{1, \ldots, n\}^n$ as type B parking functions. The number of maximal chains in $NC^B(n)$ is n^n, and Theorem 7 will show that λ labels each maximal chain with a distinct type B parking function.

Theorem 7. *The labelling λ on the type B noncrossing partition lattice is an R^*S-labelling.*

PROOF. We prove that λ is a bijection between maximal chains and sequences

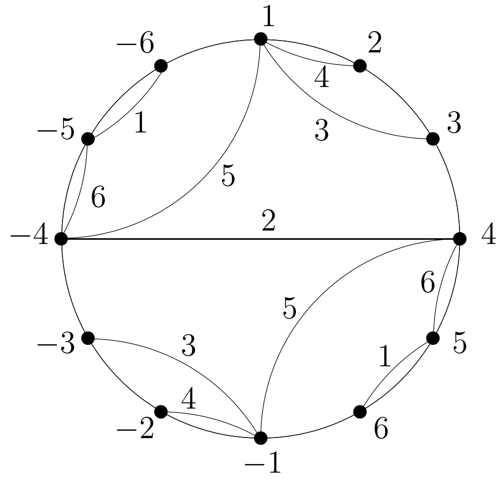

Figure 5: A maximal chain in a type B noncrossing partition lattice

in $\{1, \ldots, n\}^n$ by recursively defining the inverse map. After that, we verify that the S_n-action permuting the order of the digits in parking functions induces a local action on corresponding maximal chains to conclude that λ is an S-labelling. Finally, we check that λ is also an R^*-labelling.

It will suffice to provide the inverse to λ for those maximal chains in which 1 is merged with some $j \in \{2, \ldots, n, -1\}$ at rank 1. These chains will be labelled by exactly those parking functions which have 1 as their first digit. Symmetry will imply that for each $i \in \{1, \ldots, n\}$ there is an analogous bijection between maximal chains which have i merged with some $j \in \{i+1, \ldots, n, 1, \ldots, -i\}$ at rank 1 and type B parking functions with i as their first digit.

Given a parking function with first digit 1, the choice of which j to merge with 1 at rank 1 will depend on the content (but not the order) of the remaining $n-1$ digits in the parking function. In order to determine j, let us rearrange the remaining parking function digits in increasing order and view this weakly increasing sequence as a path from $(1, 1)$ to (n, n) made up of steps up and steps to the right, each of length 1. If every lattice point in this path is of the form (k, i) for $i \leq k$, then let $j = -1$. Otherwise, let j be the unique integer such that $(j + 1, j)$ belongs to the lattice path and this is the first place the path goes above the staircase path. (Lattice paths are used in a similar fashion in [Re].)

Of the $n-1$ digits specifying such a path, those which are less than j will determine how elements between 1 and $j-1$ will be merged in the remainder of the saturated chain while digits between j and n will determine how to merge elements between j and -1 on the circle. Note that the former correspond to type A parking functions on $1, \ldots, j-2$ while the latter correspond to elements of $[j, \ldots, n]^{n-j}$. Hence, the $n-1$ remaining digits give rise to a type A parking function on $1, \ldots, j-2$ interspersed with a type B parking function on $j, \ldots n$.

There is a corresponding decomposition of the space of maximal chains which recursively completes the bijection. This involves the following product structure which is discussed in [Re], and which generalizes similar structure for type A as studied in [SU]. Let $M_i(NC^B(n))$ denote the collection of maximal chains in the type B noncrossing partition lattice on $\pm 1, \ldots, \pm n$ which begin by merging $i \in \{1, \ldots, n\}$ with any $j \in \{i+1, \ldots, -i\}$ and merging $-i$ with $-j$. Let $M_{i,j}(NC^B(n))$ denote the restriction to maximal chains with a particular choice of j. Observe that $M_i(NC^B(n)) = \coprod_{j \in \{i+1, \ldots, -i\}} M_{i,j}(NC^B(n))$. Also note that the interval $(u, \hat{1})$ from a type B noncrossing partition u of rank 1 up to $\hat{1}$ is isomorphic to $NC^A(j-i) \times NC^B(n-j+i)$ if i is merged with j in u. Hence, maximal chains in such an interval are labelled by type A parking functions interspersed with type B parking functions, as desired.

To show that λ induces a local symmetric group action, we need to check that the maximal chain labelled by the type B parking function $(a_1, \ldots, a_n)$ differs only at rank i from the maximal chain labelled $(a_1, \ldots, a_{i+1}, a_i, \ldots, a_n)$ for $a_i \neq a_{i+1}$. Note that λ^{-1} "decides" which node to visit next at each step in choosing a maximal chain from $\hat{0}$ to $\hat{1}$ based only on the content of the remaining digits. This implies that the two maximal chains will agree up to rank $i-1$. Also observe that merge steps i and $i+1$ "commute" for $a_i \neq a_{i+1}$ by virtue of the recursively defined bijection λ. This notion of commutativity comes from treating merge steps as operators which take poset elements of rank $i-1$ to ones of rank i. Our discussion of edge-labelled graphs immediately following this theorem should clarify this point. The symmetric group relations are automatically satisfied since the action on maximal chains is induced by a valid symmetric group action on the type B parking functions which label them.

Finally, we claim that this S-labelling is also R^*. First note that the unique rising chain from $\hat{0}$ to $\hat{1}$ involves merging $\{1, \ldots, i\}$ with $\{i+1\}$ at stage i for $1 \leq i < n$ and then merging $\{1, \ldots, n\}$ with $\{-1, \ldots, -n\}$ at stage n. For $u \leq v$ the increasing chain from u to v is found similarly, but skipping steps merging components which are already merged in u or still not merged in v. $\square$

One may associate graphs to orbits and edge labellings of these graphs to maximal chains within each orbit to make the recursive structure explicit.

Figure 5 is an example of such a graph. Begin with a circle with nodes $1, \ldots, n, -1, \ldots, -n$ placed sequentially about it. For each covering relation $u \prec v$ in a maximal chain, we will draw a pair of arcs which are each labelled with the rank of v. For convenience, we will refer to left and right endpoints of arcs, by which we will mean the endpoint that appears to the left or right from the point of view from the center of the circle. The absolute value of the left endpoints of the pair of arcs labelled i will be the ith digit of the parking function which labels our maximal chain. Labelling poset edges with the absolute value of the right endpoints of the same arcs gives another poset edge-labelling. This labelling restricts to an EL-labelling for the type A noncrossing partition lattice; the type A labelling was defined by Björner and was studied by Edelman and Simion in [ES]. Unfortunately, the type B labelling is not also an EL-labelling. The type A analogue of edge-labelled graphs are equivalent to the vertex-labelled trees discussed in [ES].

If C is merged with $-C$ in the covering relation $u \prec v$, recall that there is some semicircle containing all of C. Draw an arc from i to $-i$ where i is the last element of C encountered travelling clockwise about this semicircle. If C_1 is merged with C_2 and C_1 lies entirely between consecutive elements of C_2, then draw an arc with left endpoint at the element of C_2 which comes last clockwise before reaching C_1. The right endpoint will be the first element of C_1 encountered continuing clockwise from this element of C_2. Otherwise, there is a semicircle which includes all of C_1 and all of C_2. Draw an arc connecting the nearest elements of C_1 and C_2 to each other.

These edge-labelled graphs always satisfy the following two conditions. First, each point on the circle is the right endpoint to exactly one arc. Second, the labels on the arcs with a particular left endpoint increase as one reads away from the center of the circle, i.e. as the right endpoints of these arcs progress counterclockwise.

Note that whenever two consecutive digits a_i and a_{i+1} in a parking function differ, the arcs labelled i and $i+1$ will have different left endpoints, so swapping these arc labels gives an edge-labelled graph which still satisfies the above two conditions. Hence, the two maximal chains with parking functions $(a_1, \ldots, a_n)$ and $(a_1, \ldots, a_{i+1}, a_i, \ldots, a_n)$ have edge-labelled graphs with these arc labels swapped, so the chains differ only at rank i. Successively applying adjacent transpositions shows that maximal chains p and wp in the same orbit give rise to the same underlying graph, but with arc-labels permuted by w.

One may restrict R^*S-labellings to subposets by forbidding particular arcs within these graphs. The interpolating BD noncrossing partition lattices are an example of such a restriction.

Theorem 8. *The labelling λ restricts to an R^*S-labelling for the interpolating BD noncrossing partition lattices.*

PROOF. The maximal chains of type B which do not occur in a particular interpolating BD noncrossing partition lattice constitute entire orbits. Hence the restriction of the R^*S-labelling is still an S-labelling. An increasing chain between two partitions which both satisfy $C_0 \neq \{\pm i\}$ never involves merging $\pm i$ to form $C_0 = \{\pm i\}$ at an intermediate step, so every remaining interval still has a unique increasing chain. □

Ehrenborg defined in [Eh] the function

$$F_P = \sum_{\hat{0}=t_0 \leq t_1 \leq \cdots \leq t_{k-1} < t_k = \hat{1}} x_1^{\rho(t_0,t_1)} x_2^{\rho(t_1,t_2)} \cdots x_k^{\rho(t_{k-1},t_k)}$$

which counts the chains in a finite, ranked poset P which include any particular collection of ranks. We use Theorem 8 to evaluate F_P for interpolating BD noncrossing partition lattices. Let ϕ_P be the Frobenius characteristic of the local symmetric group action induced by an R^*S-labelling of P. Simion and Stanley proved in [SS] that $F_P = \omega\phi_P$, where ω is the symmetric function involution which (in particular) swaps e_λ with h_λ.

In the following two corollaries, we sum over compositions of n, denoted by α, where α_i counts the number of digits in a type B parking function which are equal to i. Each composition accounts for a single orbit of the symmetric group action on parking functions permuting positions.

Corollary 9. *If P is the type B noncrossing partition lattice $NC^B(n)$, then*

$$F_P = \sum_{\{\alpha \in \mathbb{N}^n | \alpha_1 + \cdots + \alpha_n = n\}} e_\alpha.$$

PROOF. We invoke the theorem of Simion and Stanley that $F_P = \omega\phi_P$, with ϕ_P, as defined above. To obtain ϕ_P, we sum Frobenius characteristic over orbits. Each orbit is a product of chains and so has Frobenius characteristic h_α. □

For example, if $P = NC^B(3)$ then $F_P = e_1^3 + 6e_1e_2 + 3e_3$.

Corollary 10. *If P is the interpolating BD noncrossing partition lattice NC_S^{BD}, then*

$$F_P = \sum_{\{\alpha \in \mathbb{N}^n \cap PF_S | \alpha_1 + \cdots + \alpha_n = n\}} e_\alpha$$

for $PF_S = \bigcap_{i\in S}\{\alpha|\alpha_i+\cdots+\alpha_j < j-i+1$ *for some* $i \le j \le n$ *or* $\alpha_i+\cdots+\alpha_n+\alpha_1+\cdots+\alpha_j < j+n-i+1$ *for some* $j < i\}$.

PROOF. Theorem 8 implies $NC_S^{BD}(n)$ has an R^*S-labelling, so we again use the fact that $F_P = \omega\phi_P$. Note that the type B orbits which survive are those with edge-labelled graph not involving an arc from i to $-i$ for any $i \in S$. The definition of PF_S allows the parking functions corresponding to exactly these maximal chains. □

For example, if $P = NC_S^{BD}(3)$ with $S = \{1\}$, then $F_P = e_1^3 + 5e_1e_2 + 2e_3$. This poset has two fewer orbits than $NC^B(3)$, since the orbits of $NC^B(3)$ with label content $(1,1,1)$ and $(1,1,2)$ are no longer permitted.

In the specialization of Corollary 10 to type B, we have $S = \emptyset$, which means $\{1,\ldots,n\}^n \cap PF_S = \{1,\ldots,n\}^n$, and we recover Corollary 9. When $S = \{1,\ldots,n\}$, i.e. in the case of the type D specialization, we obtain $F_P = \sum_{a\in PF_{\{1,\ldots,n\}}} e_{\alpha(a)}$. Using recursive product structure, F_P for type D may alternatively be expressed as

$$F_{NC^B(n)} - n\left(\sum_{T\subseteq\{3,\ldots,n\}} e_{|T|+2}F_{NC^A(t_1-2)}F_{NC^A(n-t_{|T|})}\left(\prod_{i=1}^{|T|-1} F_{NC^A(t_{i+1}-t_i)}\right)\right).$$

Each of the subtracted sums comes from forbidding a particular 0-block of the form $\{\pm i\}$. The edge-labelled graphs involving an arc from i to $-i$ also have arcs from i to $i+1$, so T specifies the right endpoints of all other arcs with left endpoint i. We choose $T \subseteq \{3,\ldots,n\}$ above for the case $i = 1$, but the contribution to F_P will be the same for any i, so we multiply by n to consider all possible $i \in \{1,\ldots,n\}$. This n may be replaced by any $j \in \{1,\ldots,n\}$ to give a similar formula for an interpolating BD noncrossing partition lattice with $j = |S|$.

Acknowledgments

The author thanks Richard Stanley for suggesting this area of study, and she thanks him in addition to Clara Chan and Vic Reiner for helpful discussions and encouragement. The author is also grateful to the anonymous referee for very helpful advice about presentation.

References

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

[ES] P. Edelman and R. Simion. *Chains in the lattice of noncrossing partitons*, Discrete Math. **126** (1994), 107-119.

[Eh] R. Ehrenborg. *On posets and Hopf algebras*, Advances in Math. **119** (1996), 1-25.

[El] S. Elnitsky. *Rhombic tilings of polygons and classes of reduced words in coxeter groups*, Ph.D. thesis, Univ. of Mich., 1993.

[Fa] C.K. Fan. *Structure of a Hecke algebra quotient*, Jour. Amer. Math. Soc. **10** (1997) no. 1, 139-167.

[FR] D. Foata and J. Riordan. *Mappings of acyclic and parking functions*, aequationes math. **10** (1974), 10-22.

[Gra] D. Grabiner. *Posets in which every interval is a product of chains, and a natural local action of the symmetric group*, Discrete Math. **199**, 1-3 (1999), 77-84

[Gre] C. Greene. *Posets of shuffles*, J. Combin. Theory Ser. A **47** (1988), 191-206.

[GK] C. Greene and D. Kleitman. *Proof techniques in the theory of finite sets*, in: G.-C. Rota (ed.), Studies in combinatorics, vol. 17 of MAA Studies in Mathematics, 1978.

[Hu] J. Humphreys. Reflection Groups and Coxeter Groups, Cambridge studies in advanced mathematics 29, Cambridge University Press, Cambridge, 1990.

[Re] V. Reiner. *Non-crossing partitions for classical reflection groups*, Discrete Math. **177** (1997), 195-222.

[SS] R. Simion. and R. Stanley. *Flag-symmetry of the poset of shuffles and a local action of the symmetric group*, to appear in Discrete Math.

[SU] R. Simion. and D. Ulmann. *On the structure of the lattice of noncrossing partitions*, Discrete Math. **98** (1991), 193-206.

[St1] R. Stanley. Enumerative Combinatorics, vol. I. Wadsworth and Brooks/Cole, Pacific Grove, CA, 1986.

[St2] R. Stanley. *Flag-symmetric and locally rank-symmetric partially ordered sets*, Electronic J. Combinatorics **3 (2)**, R6 (1996), 22 pp.

[St3] R. Stanley. *Parking functions and noncrossing partitions*, Electronic J. Combinatorics **4 (2)**, R20 (1997), 17 pp.

[Ve] A.M. Vershik. *Local stationary algebras*, Amer. Math. Soc. Transl. (2) 148 (1991), 1-13; translated from Proc. First Siberian Winter School "Algebra and Analysis" (Kemerovo, 1988).