Wilf-equivalence for singleton classes

Jörgen Backelin
Department of Mathematics
University of Stockholm
Sweden

Julian West
Dept. of Mathematics and Statistics
University of Victoria
Canada

Guoce Xin
Department of Mathematics
Brandeis University
Waltham MA 02453

Abstract

Write $p_1p_2\dots p_m$ for the permutation matrix $(\delta_{p_i,j})_{m\times m}$. Let $S_n(M)$ be the set of $n\times n$ permutation matrices which do not contain the $m\times m$ permutation matrix M as a submatrix. In [3] Simion and Schmidt show bijectively that $|S_n(123)| = |S_n(213)|$. In the present work, we give a bijection from $S_n(12\dots tp_{t+1}\dots p_m)$ to $S_n(t\dots 21p_{t+1}\dots p_m)$. This result was established for $t = 2$ in [7] and for $t = 3$ in [1]. Moreover, if we think of $n\times n$ permutation matrices as transversals of the n by n square diagram, then we generalize this result to transversals of Young diagrams.

Ecrivons $p_1p_2\dots p_m$ pour indiquer la matrice (i.e. permutation) $(\delta_{p_i,j})_{m\times m}$. Soit $S_n(M)$ l'ensemble des matrices permutationelles $n\times n$ qui ne contient pas une sous-matrice $m\times m$ isomorphe á une matrice (permutation) M fixe. Simion et Schmidt ([3]) ont montré bijectivement que $|S_n(123)| = |S_n(213)|$. Nous presentons un bijection entre $S_n(12\dots tp_{t+1}\dots p_m)$ et $S_n(t\dots 21p_{t+1}\dots p_m)$. Ce résultat à été établi pour $t = 2$ dans [7] et pour $t = 3$ dans [1]. Nous pensons des matrices permutationelles comme traverses du diagram carrée n par n, et generalisons ce résultat aux traverses des diagrammes de Young.

1 Introduction

The systematic study of permutations avoiding forbidden subsequences was launched in 1985 by the seminal paper of Rodica Simion and Frank Schmidt [3]. In this paper we give a best-possible generalisation of one of the proofs from that paper.

We begin with some necessary definitions, which we will cast in the form of statements about permutation matrices. The matrix notation adds a degree of symmetry to the problem, which will be exploited in our proofs.

A *permutation matrix of order n* is a transversal of the n by n square diagram, in other words, a placement of n non-attacking rooks on an n by n board.

Given a permutation matrix M of order m, a permutation matrix N of order $n > m$ will be said to *contain* the smaller matrix M if there exist two subsets of the index set $[n]$, $R = \{r_1 < r_2 < \cdots < r_m\}$ and $C = \{c_1 < c_2 < \cdots < c_m\}$, such that

$$\begin{bmatrix} N[r_1, c_1] & N[r_1, c_2] & \cdots & N[r_1, c_m] \\ N[r_2, c_1] & N[r_2, c_2] & \cdots & N[r_2, c_m] \\ \vdots & \vdots & & \vdots \\ N[r_m, c_1] & N[r_m, c_2] & \cdots & N[r_m, c_m] \end{bmatrix} = M. \tag{1.1}$$

If a permutation matrix N does not contain a given submatrix M, then it is said to *avoid* M. Let $S_n(M_1, M_2, \ldots, M_q)$ be the set of permutation matrices of order n which avoid each of $M_1, M_2, \ldots, M_q$. This is simply a restatement in terms of matrices of the theory of *permutations with forbidden subsequences*, which has been developed in [3, 4, 5, 7, 8, 9]. In order to complete the translation to this language, we abbreviate by $p_1p_2 \ldots p_m$ the matrix $M_{i,j} = (\delta_{p_i,j})$.

The Simion-Schmidt paper of 1985 produces a comprehensive catalogue of results for restrictions when each M_i is of order 3. Among these is a new proof of the previously known fact that the patterns 123 and 213 are equally restrictive. In fact, $|S_n(123)| = |S_n(213)| = \frac{1}{n+1}\binom{2n}{n}$, the n-th Catalan number.

The problem of identifying equirestrictive sets of forbidden patterns was posed by Herb Wilf in the early 1980s. Therefore we will say, following the practice of [1] that two sets of matrices $\mathcal{M} = \{M_1, M_2, \ldots, M_q\}$ and $\mathcal{M}' = \{M_1', M_2', \ldots, M_{q'}'\}$ are *Wilf-equivalent* if $|S_n(M_1, M_2, \ldots, M_q)| = |S_n(M_1', M_2', \ldots, M_{q'}')|$ for all positive integers n. It is almost obvious that $\mathcal{M}$ and $\mathcal{M}'$ are Wilf-equivalent if $q = q'$ and each M_i' can be obtained from the corresponding M_i by an operation of a fixed element of the dihedral group D_4; in this case the elements of $S_n(M_1', M_2', \ldots, M_{q'}')$ are just the elements of $S_n(M_1, M_2, \ldots, M_q)$ after application of the same operation.

Obtaining any sort of general results about Wilf-equivalence, or related questions, appears to be quite hard. Attention has focused on establishing special cases. Generally it is taken for granted that $q = q'$ and that each M_i has the same order as M_i'. These are natural conditions, but it has never been demonstrated that they are (in any non-trivial sense) necessary. They are also far from sufficient.

Somewhat oddly, it seems that Wilf-equivalence is a more widespread phenomenon for sets $\mathcal{M}$ and $\mathcal{M}'$ with $q = q' > 1$ than for singleton classes. For instance $\{3142,\ 2413\}$ is Wilf-equivalent to $\{2431,\ 4231\}$ (see [8]), and there are numerous further examples in [9]. Indeed, the purpose of the present work is to establish a result which takes care of almost all cases where $\mathcal{M} = \{M\}$ and a Wilf-equivalence can be conjectured after computer-assisted examination of small cases. In all these cases, $\mathcal{M}' = \{M'\}$, M and M' are of the same order m, and

$$M = \begin{bmatrix} I_t & 0 \\ 0 & A \end{bmatrix}, \quad M' = \begin{bmatrix} J_t & 0 \\ 0 & A \end{bmatrix}, \tag{1.2}$$

where $I_t = 12 \cdots t$ and $J_t = t \cdots 21$.

The only other pair of matrices known to be Wilf-equivalent is the following (for a proof see [5]):

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix}, \; M' = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \tag{1.3}$$

There is only one more possible pair for $q = q' = 1$ and $m < 7$ other than those related by (1.2) and (1.3), possibly after the application of transformations from the dihedral group as explained above. This pair, 231564 and 231645, was first observed by Zvezdelina Stankova-Frenkel (personal communication). A proof of the Wilf-equivalence of these two permutations will appear in a paper by Stankova and West [6].

In [3], Wilf-equivalence was established for the two matrices which correspond to (1.2) where $t = 2$ and $A = I_1$. In [7] this was extended to $t = 2$ and arbitrary A, while in [1] the cases of both $t = 2$ and $t = 3$ were established. The computer search for $m < 7$ suggests that (1.2) is the best possible general result for singleton classes. (The existence of further sporadic cases can not, of course, be ruled out.) In this paper we prove the general case of arbitrary t, simultaneously extending the result to cover Young diagrams.

Each stage in the development of this proof through [3], [7] and [1] to the present paper relied heavily on its predecessor. Thus it is reasonable to assert that this paper represents a direct generalisation of Simion and Schmidt's original proof.

2 The Main Theorem

We will actually be proving an extended version of the conjecture that the matrices of equation 1.2 are Wilf-equivalent. We extend the definition of a forbidden submatrix from an n by n permutation matrix N to a transversal of a Young diagram $\lambda = (\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n)$. A transversal L of the shape λ is an assignment of zeroes and ones to the squares of λ, such that each row and column contains exactly one 1. We note that such a transveral can only exist if, at the very least, $\lambda_1 = n$. Since a permutation matrix is simply a transversal of the square shape $\lambda_1 = \cdots = \lambda_n = n$, transversals of Young diagrams include permutation matrices as a special case.

The transversal L will be said to *contain* the matrix M if there exist two subsets of the positive integers $R = \{r_1 < r_2 < \cdots < r_m\}$ and $C = \{c_1 < c_2 < \cdots < c_m\}$, such that *each of the squares* (r_i, c_j) *falls within the board*, and the matrix equation (1.1) holds. Since λ is a Young diagram, if (r, c) falls in λ then the rectangle with corners $(1, 1)$ and (r, c) will be contained in λ, so we need only that (r_m, c_m) falls within the board. Let $S_\lambda(M)$ be the set of all transversals that avoid M.

As we are regarding the shapes being transversed as having the matrix shape $\lambda_1 = \ldots \lambda_n = n$ as a special case, it is natural that we should use the matrix conventions for describing Young diagrams. That is, the row index increases from top to bottom and the column index from left to right, so that we may speak of a cell (r', c') as being *above* and *to the right* of (r, c) if $r' < r$ and $c' > c$, etc.

Definitions and notation. Recall that two sets of permutation matrices $\mathcal{M}$ and $\mathcal{M}'$ are said to be Wilf-equivalent if $|S_n(\mathcal{M})| = |S_n(\mathcal{M}')|$ for all natural numbers n. We will use the notation $\mathcal{M} \underset{W}{\sim} \mathcal{M}'$ to indicate that $\mathcal{M}$ and $\mathcal{M}'$ are Wilf-equivalent. In this paper we are concerned with examples where $\mathcal{M}$ and $\mathcal{M}'$ each consist of a single matrix; we can use the notation $M \underset{W}{\sim} M'$ for such singleton classes.

For our extended notion of forbidden patterns, we will say that two permutation matrices M and M' are *shape-Wilf-equivalent* if $|S_\lambda(M)| = |S_\lambda(M')|$ for all partitions λ. We denote shape-Wilf-equivalence by $M \underset{sW}{\sim} M'$ (or $\mathcal{M} \underset{sW}{\sim} \mathcal{M}'$). ∎

We can now present the main theorem.

Theorem 2.1 *Let M and M' be related by*

$$M = \begin{bmatrix} I_t & 0 \\ 0 & A \end{bmatrix}, \; M' = \begin{bmatrix} J_t & 0 \\ 0 & A \end{bmatrix} \tag{2.1}$$

Then $M \underset{sW}{\sim} M'$.

It is easy to see that the above theorem is a consequence of the following two propositions.

Proposition 2.2 *For all $t > 0$, $I_t \underset{sW}{\sim} J_t$.*

Proposition 2.3 *Let C, D be two matrices of order t, and let M and M' be related by*

$$M = \begin{bmatrix} C & 0 \\ 0 & A \end{bmatrix}, \; M' = \begin{bmatrix} D & 0 \\ 0 & A \end{bmatrix} \tag{2.2}$$

If $C \underset{sW}{\sim} D$ then $M \underset{sW}{\sim} M'$.

Proof of Proposition 2.3. Let M and M' be related as in (2.2). For every Young diagram π, Let Π_π be a bijection from $S_\pi(C)$ to $S_\pi(D)$. (The existence of such bijections is guaranteed by the hypothesized shape-Wilf-equivalence of C and D.) We construct a bijection α from $S_\lambda(M)$ to $S_\lambda(M')$ as follows.

Suppose $N \in S_\lambda(M)$.

Step 1: For any square $(i, j) \in \lambda$, if the subboard of λ that is below and to the right of (i, j) contains A, then colour it white; otherwise, colour it blue.

Step 2: Find the 1's coloured blue, and colour the corresponding rows and columns blue.

Step 3: Denote the white board by π, and the transversal on π by L. Transform L to $\Pi_\pi(L)$. Then we get an $\alpha(N) \in S_\lambda(M')$.

In order to apply Π_π in step 3, it is necessary that π be a Young diagram. To see that it is, first note that if any square (r, c) is coloured white in step 1, then so are all squares in the rectangle with corners $(1, 1)$ and (r, c). The squares which are white after step 1 therefore form a Young diagram, and the 1s in it form a partial transversal. In step 2, some rows and columns of this diagram are deleted, retaining only those which contain the 1's of the partial transversal; the result of this procedure is still a Young diagram, which we have called π, and the 1's in it now form an actual transversal, L.

The inverse procedure will be identical except for replacing Π_π by Π_π^{-1} in step 3. Suppose $N' \in S_\lambda(M')$.

Step 1: For any square $(i, j) \in \lambda$, if the subboard of λ below and to the right of (i, j) contains A, then colour it white; otherwise, colour it blue.

Step 2: Find the 1's coloured blue, and colour the corresponding rows and columns blue.

Step 3: Denoted the white board by π', and the transversal on π' by L', transform L' to $\Pi_\pi^{-1}(L')$. Then we get an $\alpha^{-1}(N') \in S_\lambda(M)$.

Since both Π_π and Π_π^{-1} change only the 1's in the white board, the inverse procedure gets the same colours for any squares of λ, i.e., $\pi = \pi'$.

Notice that $N \in S_\lambda(M)$ if and only if $L \in S_\pi(C)$, and $N' \in S_\lambda(M')$ if and only if $L' \in S_\pi(D)$. Then theorem 2.3 follows. ∎

It is easy to see that we can replace the A in the above proposition by a class of matrices. The proof requires no more ideas.

Note that proposition 2.3 does not apply to ordinary Wilf-equivalence. For instance, $1234 \underset{W}{\sim} 2143$, but 123456 is not Wilf-equivalent to 214356, as $|S_9(123456)| = 344,837$ but $|S_9(214356)| = 344,838$.

In the remainder of this paper, we offer two proofs of proposition 2.2. In [6], it is further proved that $(231) \underset{sW}{\sim} (312)$. There are no further cases where shape-Wilf-equivalence is known or conjectured; moreover all known or conjectured cases of Wilf-equivalence are explained by considerations discussed above.

3 The First Proof (Strategy)

Instead of proving Proposition 2.2, we will state and prove an equivalent Proposition 3.1. We will give the proof of Proposition 3.1 in this and the following section.

First, for all natural numbers t, define the matrix F_t as follows:

$$F_t = \begin{bmatrix} J_{t-1} & 0 \\ 0 & 1 \end{bmatrix}. \tag{3.1}$$

Proposition 3.1 *For all $t > 0$, $F_t \underset{sW}{\sim} J_t$.*

Proof of the equivalence of propositions 2.2 and 3.1. Suppose proposition 2.2 is true. Then $J_t \underset{sW}{\sim} I_t$. Furthermore, $I_{t-1} \underset{sW}{\sim} J_{t-1}$, so it follows that $I_t \underset{sW}{\sim} F_t$, using proposition 2.3 with $C = I_{t-1}$, $D = J_{t-1}$ and $A = I_1$. Proposition 3.1 follows immediately.

Now suppose proposition 3.1 is true. Let $G_{t,k}$ be defined by

$$G_t = \begin{bmatrix} J_{t-k} & 0 \\ 0 & I_k \end{bmatrix}, \tag{3.2}$$

so that $G_{t,0} = J_t$, $G_{t,1} = F_t$ and $G_{t,t} = I_t$. For all $0 \leq k < t$, proposition 2.3 using $C = J_{t-k}$, $D = F_{t-k}$ (which are shape-Wilf-equivalent using proposition 3.1) and $A = I_k$ shows that $G_{t,k} \underset{sW}{\sim} G_{t,k+1}$. Iterating this process establishes $G_{t,0} \underset{sW}{\sim} G_{t,t}$, which is proposition 2.2. ∎

We will prove proposition 3.1 by exhibiting a bijection between $S_\lambda(F_t)$ and $S_\lambda(J_t)$, valid for all shapes λ and for all natural numbers t. The key step in the bijection from F_t-avoiding transversals into the class of J_t-avoiding transversals will be to identify canonical examples of the forbidden pattern J_t and replace them with copies of F_t.

Let G be a submatrix isomorphic to J_t at rows $r_1 < r_2 < \cdots < r_t$, and columns $c_1 < c_2 < \cdots < c_t$. Then the squares (r_i, c_{t-i+1}) are filled by 1's for $i = 1, 2, \ldots, t$, and all the other squares are filled by 0's.

Let $\theta(G)$ be the submatrix isomorphic to F_t with the same squares as G. Now we can give our main bijection as follows.

Bijection from $S_\lambda(F_t)$ to $S_\lambda(J_t)$.

Suppose L is a transversal in $S_\lambda(F_t)$.

Step 1: If L contains no J_t, end.

Step 2: Find the highest square a_1 containing a 1 such that there is a J_t in L, in which a_1 is the left most 1.

Step 3: Find the leftmost square a_2 containing a 1 such that there is a J_t in L, in which a_1 and a_2 are the left most two 1's.

Step 4: Find $a_3, a_4, \ldots, a_t$ one by one as in Step 3. Then we get a G, which is isomorphic to J_t.

Step 5: Replace G by $\theta(G)$ and leave all the other squares fixed.

Step 6: Repeat the above procedure until there is no J_t left. Then we get a $T \in S_\lambda(J_t)$.

Remark: If L contains J_t, we denote the result of applying steps 2–5 to L by $\phi(L)$.

Bijection from $S_\lambda(J_t)$ to $S_\lambda(F_t)$.

To construct the inverse bijection, suppose T is a transversal in $S_\lambda(J_t)$.

Step 1: If T contains no F_t, end.

Step 2: Find the lowest b_t containing a 1 such that T contains an F_t, in which b_t is the rightmost 1.

Step 3: Find the lowest b_{t-1} containing a 1 such that T contains an F_t, in which b_t, b_{t-1} are the rightmost two 1s.

Step 4: Find $b_{t-2}, b_{t-3}, \ldots, b_1$ one by one as b_{t-1}. Then we get an H, which is isomorphic to F_t.

Step 5: Replace H by $\theta^{-1}(H)$ and leave all the other squares fixed.

Step 6: Repeat the above procedure until there is no F_t left. Then we get an $L \in S_\lambda(F_t)$.

Remark: If T contains F_t, we denote the result of applying steps 2–5 to T by $\psi(T)$.

In the next section, we show that the above two procedures really are inverse to each other. This will follow because the two functions ϕ and ψ are inverse to each other in appropriate cases. This will be proved by six lemmas in the next section. Since our main bijections are just produced by iterating ϕ (ψ) until there are no remaining examples of J_t (F_t), Proposition 3.1 will follow.

4 The First Proof (Correctness of the Bijection)

To prove the correctness of our main bijection, we need to show the following two facts about ϕ. One is that our algorithm terminates, which is confirmed by Lemma 4.1. It says that the a_1 we find in the algorithm ϕ can only go down or slide right. The other one is that when we apply the inverse algorithm ψ, b_i will not run away for $i = 1, 2, \ldots, t$. This is confirmed by Lemma 4.2 and Lemma 4.3.

Also there are two facts about ψ we must show. One is that the inverse procedure of our algorithm terminates, which is confirmed by Lemma 4.4. It says that the b_t we find in the algorithm ψ can only go up or slide left. The other one is that when we apply the inverse algorithm ϕ, a_i will not run away for $i = 1, 2, \ldots, t$. This is confirmed by Lemma 4.5 and Lemma 4.6.

Using the notation of the algorithm for ϕ, label the t squares which contain 1's in $\theta(G)$, from left to right, by $b_1, b_2, \ldots, b_t$. Denote by E the union of the following three parts of the board: the board that is above and to the right of a_{t-1} but to the left of a_t, the board that is above and to the left of a_1 but below b_1, and the union of the rectangles with corners a_i and b_{i+1} for $i = 1, 2, \ldots, t-2$. (see the following figure).

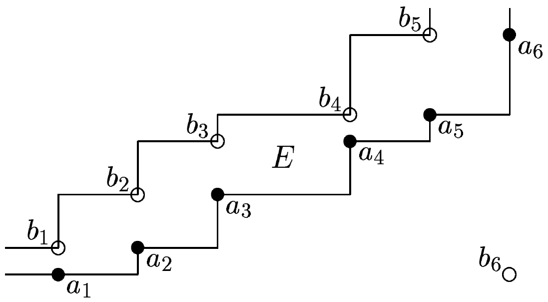

Figure 1: The region E (for $t = 6$)

It is clear that E separates λ into two parts, a part to the left of (and above) E and a part to the right of (and below) E. Moreover, all the 1's of λ are in one or the other of these two parts, as there are no 1's inside E by the selection of a_i's. For instance, if there were a 1 in the rectangle with corners a_i and b_{i+1}, then that 1 would have been selected in preference to a_{i+1}, as it lies to the left of a_{i+1} and forms a J_t together with the rest of the a_i's. The board E plays a leading role, because its shape makes us go at least "two steps" whenever we ascend across it.

All our proofs are by contradiction. The main idea is to replace some 1's of an assumed F_t or J_t by a_s or b_s ($s = 1, 2, 3, \ldots, t$) to form an F_t or J_t which contradicts the way in which the a_i's and b_i's were selected. First of all, these substitutions should preserve the relative positions of the 1's, and should not decrease the number of the 1's. Since it is not true that any t ascending 1's will form a J_t in λ (the cell whose row is the that of the lowest 1 and whose column that of the rightmost 1 must also lie within the board), we should also keep in mind that the lower right corner of the assumed F_t or J_t should not go right or lower.

We need the following three useful properties that provide bounds on the longest anti-identities which can lie in certain subregions of the area to the left of E.

Properties.

(1) For any $1 \le i \le t-1$, the board that is above a_1 and below b_i can't contain a J_i with all its 1's to the left of E in L (or $\phi(L)$).

(2) For any $1 \le i \le t-1$, the board that is above and to the right of b_i can't contain a J_{t-i} with all its 1's to the left of E in L (or $\phi(L)$).

(3) For any $1 \le i < j \le t-1$, the rectangle with corners b_i and b_j can't contain a J_{j-i} with all its 1's to the left of E in L (or $\phi(L)$).

Proof. We prove the three properties for L.

(1) If there is such a J_i below b_i, then it is below a_{i+1}. Therefore these i 1's, combining with $a_{i+1}, a_{i+2}, \ldots, a_t$, will form a J_t in L, which contradicts the selection of a_1.

(2) If there is a J_{t-i} in this region, then either its leftmost 1 is to the left of b_{i+1} (and hence to the left of a_{i+1}), or else it all lies to the right of b_{i+1} (and a_{i+1}). In the first case, this J_{t-i}, combining with $a_1, a_2, \ldots, a_i$, will form a J_t in L, which contradicts to the selection of a_{i+1}. In the second case, $a_2, a_3, \ldots, a_{i+1}$, combining with this J_{t-i}, will form a J_t in L, which contradicts the selection of a_1.

(3) If there is a J_{j-i} in this region, then either its leftmost 1 is to the left of a_{i+1}, or else it all lies to the right of a_{i+1}. In the first case, so $a_1, a_2, \ldots, a_i$, combining with this J_{j-i} and $a_{j+1}, \ldots, a_t$, will form a J_t in L, which contradicts to the selection of a_{i+1}. In the second case, $a_2, a_3, \ldots, a_{i+1}$, combining with this J_{j-i} and $a_{j+1}, \ldots, a_t$ will form a J_t in L , which contradicts the selection of a_1. ∎

Remark: In each of these subregions of the area left of E, the largest possible number of ascending 1's in the region is just the number of b_i's in the corresponding part of the board.

Lemma 4.1 *The rows above a_1 can't contain a J_t in $\phi(L)$.*

Proof of 4.1 If not, suppose G is such a J_t in $\phi(L)$. Label its 1's by $g_1, g_2, \ldots, g_t$ from left to right. Since the algorithm does not change the other 1's positions, at least one of $b_1, b_2, \ldots, b_t$ must fall in G. We shall replace some 1's of G to form a J_t in L. This will contradict the selection of a_1.

Find the smallest i such that b_i falls into G. Then apply the following algorithm.

Horizontal slide algorithm for ϕ: If there is a 1 of G that is above b_i and to the right of E, find the first 1 and denote it by g_y. Find x such that g_y is below b_x, and above b_{x-1}. Then by property (3), there are fewer than $x-i$ 1's in G that are below b_x, but not below b_i and to the left of E. So we can replace these 1's by $b_i, b_{i+1}, \ldots, b_{x-1}$, and hence by $a_{i+1}, a_{i+2}, \ldots, a_x$, which are in L.

We can repeat this horizontal slide algorithm until one of the following two cases appears.

1. There is no b_i that falls in G. This ends the proof.

2. There is such a b_i, but g_y is above a_t, or g_y does not exist. By property (2), there can't exist $t-i+1$ 1's of G which are to the right of a_i and to the left of E. So we can replace them by $a_{i+1}, a_{i+2}, \ldots, a_t$. Then we have a J_t in L that is above a_1 as desired. ∎

Lemma 4.2 *If L contains no F_t with at least one square in a row below a_1, then $\phi(L)$ contains no such F_t.*

Proof of 4.2 If not, suppose H is such an F_t in $\phi(L)$. Label the t 1's in H by $h_1, h_2, \ldots, h_t$ from left to right. Then h_t is below a_1. We shall replace some 1's of H (except h_t) to form an F_t in L, which will contradict the hypothesis.

By the hypothesis, h_t can't form an F_t in L, so h_t must be at the left side of a_{t-1}. Notice that the algorithm ϕ only changes a_i to b_i for $i = 1, 2, \ldots, t$. At least one of b_i must fall in H.

Find the largest i such that b_i falls into H.

Vertical slide algorithm for ϕ: If there is a 1 of H which is below b_i and to the right of E, find the first 1 and denote it by h_y. Find x such that h_y is to the right of b_x, and to the left of b_{x+1}. By property (3), there are at most $i - x$ 1's in H, which are above b_x but not above b_i, and to the left of E. So we can replace them by $b_{x+1}, b_{x+2}, \ldots, b_i$, and hence by $a_{x+1}, a_{x+2}, \ldots, a_i$.

We can repeat the vertical slide algorithm until the following two cases appear.

1. There is no b_i that falls in H. This ends the proof.

2. There is a such b_i, but h_y does not exist. By property (1), there are at most i 1's of H that are not above b_i and to the left of E. So we can replace them by $a_1, a_2, \ldots, a_i$ to form an F_t in L. Since we never touch h_t, the lemma follows. ∎

Lemma 4.3 *The board that is to the left of a_{k+1} and above a_1 can't contain a J_k in $\phi(L)$ with its highest 1 below b_k.*

Proof of 4.3 If not, suppose G is such a J_k in $\phi(L)$. Label its 1's by $g_1, g_2, \ldots, g_k$ from left to right. Since the algorithm does not change the other 1's positions, one of $b_1, b_2, \ldots, b_k$ must fall in G. Otherwise, these k 1's, combining with $a_{k+1}, a_{k+2}, \ldots, a_t$, will form a J_t in L. This contradicts the selection of a_1.

Find the smallest i such that b_i falls in G. If g_k is to the right of E, we can repeat the Vertical slide algorithm unti l all these k 1's are replaced to fall in L. Then these k 1's, combining with $a_{k+1}, a_{k+2}, \ldots, a_t$, will form a J_t in L. This contradicts the selection of a_1. Otherwise, we can repeat the horizontal slide algorithm, and get a J_k in $\phi(L)$ with all its 1's below b_k and to the left of E, this contradicts property (1). ∎

With all the notation as in the algorithm ψ, the diagram E as in Figure 1 also contains no 1's when there is no J_t above a_1 in T. It will play a leading role for the same reason.

There is a similar property about the largest number of 1's to the right of E.

Property.

(1′) For any $1 \leq i < j \leq t$, the rectangle with corners a_i and a_j can't contain a J_{j-i} with all its 1's to the right of E in T or $\psi(T)$.

Proof. We prove the property for T. The property for $\psi(T)$ is similar.

If not, $b_1, b_2, \ldots, b_{i-1}$ (when i=1, this means empty squares), combining with this J_{j-i} and $b_j, b_{j+1}, \ldots, b_t$, will form an F_t in T. This contradicts the selection of b_{j-1}. ∎

Remark: The largest number of ascending 1's to the right of E is just the number of a_i's in the corresponding board.

Lemma 4.4 *If T contains no F_t with at least one square in a row below a_1, then $\psi(T)$ contains no such F_t.*

Proof of 4.4 If not, suppose H is such an F_t in $\psi(T)$. Label the t 1's in H by $h_1, h_2, \ldots, h_t$ from left to right. Then h_t is below a_1. Similar as in the proof of Lemma 4.2, we shall replace some 1's of H (except h_t) to form an F_t in T.

By the selection of b_t, h_t can't form an F_t in T, so h_t must be at the left side of b_{t-1}. Notice that the algorithm ψ only changes b_i to a_i for $i = 1, 2, \ldots, t$. At least one of the a_i's must fall in H.

Find the smallest i such that a_i falls into H.

Vertical slide algorithm for ψ: If there is a 1 of H which is above a_i and to the left of E, find the f irst 1 and denote it by h_y. Find x such that h_y is to the right of a_x and to the left of a_{x+1}. Then by property $(1')$, there are at most $x - i + 1$ 1's in H that are below a_{x+1}, not below a_i, and to the right of E. So we can replace these 1's by $a_i, a_{i+1}, \ldots, a_x$, and hence by $b_i, b_{i+1}, \ldots, b_x$, which are in T.

We can repeat the vertical slide algorithm until one of the following two cases appears.

1. There is no a_i that falls in H. This ends the proof.

2. There is a such a_i, but h_y does not exist. Then suppose a_v is the first 1 to the right of h_t. By property $(1')$, there are at most $v - i$ 1's of H that is below and to the left of a_v, but not below a_i, and to the right of E. So we can replace these 1's by $a_i, a_{i+1}, \ldots, a_{v-1}$, and hence by $b_i, b_{i+1}, \ldots, b_{v-1}$. Then we have an F_t in T with a square h_t below a_1. ∎

Lemma 4.5 *If T contains no J_t that is above a_1, neither does $\psi(T)$.*

Proof of 4.5 If not, suppose G is such a J_t in $\psi(T)$. Label its 1's by $g_1, g_2, \ldots, g_t$ as before. Since the algorithm ψ does not change the other 1's positions, at least one of $a_1, a_2, \ldots, a_t$ must fall in G.

We prove the lemma by replacing some 1's of G to form a J_t in T that is above a_1. This contradicts the hypothesis.

Find the largest i such that a_i falls into G.

Horizontal slide algorithm for ψ: If there is a 1 of G which is below a_i and to the left of E, find the rightmost 1 and denote it by g_y. Find x such that g_y is below a_x, and above a_{x-1}. Then by property $(1')$, there are at most $i - x + 1$ 1's in G that are above a_{x-1}, not above a_{i+1}, and to the right of E. So we can replace these 1's by $a_x, a_{x+1}, \ldots, a_i$, and hence by $b_{x-1}, b_x, \ldots, b_{i-1}$.

We can repeat this horizontal slide algorithm until one of the following two cases appears.

1. There is no a_i that falls in G. This ends the proof.

2. There is such an a_i, but g_y does not exist. Find x such that g_1 is below a_{x+1} and above a_x. Then by property $(1')$, the number of the 1's of G that are above a_x but not above a_i are no more than $i - x$, so we ca n replace these 1's by $a_{x+1}, a_{x+2}, \ldots, a_i$, and hence by $b_x, b_{x+1}, \ldots, b_{i-1}$. Then we get a J_t in T which is above a_1. ∎

Lemma 4.6 *If T contains no J_t that is above a_1, the board that is above and to the right of a_k can't contain a J_{t-k} in $\psi(T)$ such that the lowest 1 of this J_{t-k} is to the left of a_{k+1}, and this J_{t-k}, combining with $a_1, a_2, \ldots, a_k$, forms a J_t in $\psi(T)$.*

Proof of 4.6 If not, suppose G is such a J_{t-k} in $\psi(T)$. Label its 1's by $g_{k+1}, g_{k+2}, \ldots, g_t$ from left to right. By hypothesis, g_{k+1} is to the left of a_{k+1}. Since the algorithm ψ does not change the other 1's positions, one of $a_{k+1}, a_{k+2}, \ldots, a_t$ must fall in G. Otherwise, this J_{t-k}, combining with $b_1, b_2, \ldots, b_k$, will form a J_t in T which is above a_1.

Find the largest i such that a_i falls into G.

Because of the position of g_{k+1}, we can repeat the horizontal slide algorithm for ψ until no a_i's exist. Then we get a J_{t-k} in T. This J_{t-k}, combining with $b_1, b_2, \ldots, b_k$, will form a J_t in T which is above a_1 . ∎

We can now assemble the proof of proposition 3.1, which is the last step in the proof of 2.1. Conceptually, the proof works as follows. The operation of the forward algorithm created by iterating ϕ is to locate J_t's and convert them into F_t's. This is done from the top down. Therefore, at any time during the operation of this algorithm, there is an "F" region at the top of the matrix (Young diagram) containing F_t's and a "J" region at the bottom containing J_t's. The function ϕ locates the highest J_t in the lower region and converts it to an F_t, pushing it across the boundary into the upper region. The inverse function ψ pushes it back again.

Of course in an arbitrary matrix containing a mix of both J_t's and F_t's, applying ϕ will convert the highest J_t, which does not necessarily become the lowest F_t identified by ψ, so the functions will not necessarily invert one another. However, since we start out with a matrix containing only J_t's and no F_t's, we know we will at all times have an upper F region and a lower J region. Within this context, the functions ϕ and ψ are inverses, and therefore their iterates will carry us back and forth from a matrix which is entirely J region to one which is entirely F region. The map from $S_\lambda(F_t)$ to $S_\lambda(J_t)$ is therefore bijective.

Proof of proposition 3.1 Consider the bijection from $S_\lambda(F_t)$ to $S_\lambda(J_t)$. Suppose we start with some $L \in S_\lambda(F_t)$. At the n-th application of ϕ we select a copy of J_t in $\phi^{n-1}(L)$. This has its lowermost 1 in some row r. By lemma 4.1, the J_t we will select in $\phi^n(L)$ cannot have its lowermost 1 anywhere above row r. If it is in row r, then we know it is further to the right than at the previous iteration, because there is only one 1 in that row, and we have just moved it to the right, form a_1 to b_t. It follows that at each iteration the selection of a_1 proceeds lexicographically through the diagram λ, and therefore the algorithm must terminate.

Next we must show that the F_t identified by ψ in $\phi^n(L)$ will be the one just created by the application of ϕ to $\phi^{n-1}(L)$. The function ψ is defined by an algorithm, which selects the 1's of ψ sequentially. The first one to be selected is b_t; this will by definition be the lowest 1 in the matrix which forms a b_t in some F_t. We know that the 1 in row r has this property because we just created it in the previous application of ϕ. We also know by lemma 4.2 that it is the lowest 1 to have this property; it will therefore be the one selected. (It is clear that the conclusion of lemma 4.2 guarantees just what we want; the hypothesis of lemma 4.2 is valid by induction because $\phi^n(L)$ is simply created by repeated application of ϕ to L, a transversal which initially contained no F_t's at all.)

Next we must select a 1 so that the two 1's selected so far form the rightmost entries of some F_t. Again, the 1 we have just positioned at b_{t-1} has this property. We are futhermore looking for the lowest such 1; we claim that there can be no suitable 1 lower than b_{t-1} by the application of lemma 4.3 with $k = t - 1$. Therefore b_{t-1} is indeed selected. The remaining selections of $b_{t-2}, \ldots, b_1$ are likewise guaranteed by the application of lemma 4.3 with $k = t-2, \ldots, 1$.

This completes the proof that $\psi(\phi(\phi^{n-1}(L))) = \phi^{n-1}(L)$. If it takes N applications of ϕ to remove *all* the J_t's from L (i.e. $\phi^N(L) \in S_\lambda(J_t)$), then we can conclude that $\psi^N(\phi^N(L)) = L$, in other words that the two bijections of section 3 are inverse in this sense.

The proof that equally $\phi^{N'}(\psi^{N'}(T)) = T$ is essentially the same, relying on the lemmata 4.4, 4.5 and 4.6. ∎

5 The Second Proof (Strategy)

In the second proof we will need to refer to the *outside corners* of λ, which are the maximal elements of λ viewed as a finite order ideal of $\mathbf{N} \times \mathbf{N}$. The outside corners form an antichain in this order ideal, and can themselves be viewed as ordered (top to bottom) by letting $(r_1, c_1) \leq' (r_2, c_2)$ iff $r_1 \leq r_2$ but $c_1 \geq c_2$.

Let $R(a,b,c,d) := \{(i,j) \in \mathbf{N} \times \mathbf{N} : a \leq i \leq b \text{ and } c \leq j \leq d\}$ be called a *rectangle* and $R'(b,d) := R(0,b,0,d)$ be called a *shadow*. Then by a *shaded Young diagram* we will mean a pair (λ, R) of a shape λ and a shadow R, such that either $R = \emptyset$ or R reaches the outer border of λ, in the sense that $R = R'(r,c)$ with $(r+1, c+1) \notin \lambda$. The *shaded part* of *lambda* is $\lambda \cap R$. Unless $R = \emptyset$, the *unshaded part* $\lambda \setminus R$ is divided into two parts, those elements strictly below R and those strictly to the right of R, respectively. If E is a set of squares, then $R'(E)$ will denote the minimal shadow containing all elements of E.

We will be concerned not merely with transversals but with partial transversals (PT), in which each row and column contains at most one 1.

For any $p = (r,c)$ we may form the *contraction* λ/p of λ with respect to p by removing p together with all other squares in its row and its column, and joining up the remainder. The *expansion* μ of λ with respect to p is the minimal set such that $\mu/p = \lambda$.

Note that a contraction of a Young diagram, rectangle, shadow or PT is itself of the same type. Also, if (λ, R) is a shaded pair, then so is $(\lambda/p, R/p)$. Moreover, if (λ, R) is a shaded Young diagram and $p = (a,b)$ is an outside corner of λ, and either $R = \emptyset$ or p is one of the outside corners adjacent to R (that is, $p \in \lambda \setminus R$ and there is no other outside corner $q \in \lambda \setminus R$ between p and R), then $(\lambda \cap R) \cup R'(a,b) = \lambda \cap (R'(R \cup \{p\}))$ and $(\lambda, R'(R \cup \{p\}))$ is a shaded Young diagram.

Extending the notation of the previous sections, we will continue to say that a (partial) transversal L contain a permutation M if 1.1 can be satisfied with $(r_m, c_m) \in \lambda$, and for $q \geq 0$, $S^q_\lambda(M) := \{PTL \text{ of } \lambda : |L| = q \text{ and } L \text{ avoids } M\}$.

We wish to prove proposition 2.2, which states that $|S_\lambda(I_t)| = |S_\lambda(J_t)|$ for all shapes λ and $t > 0$. We immediately generalise this to

Proposition 5.1 $|S^q_\lambda(I_t)| = |S^q_\lambda(J_t)|$ *for all shapes* λ*,* $t > 0$*,* $q \geq 0$*.*

We will generalise still further by making some extended definitions for $\lambda = I_t, J_t$. Let $S^q_{\lambda,R}(I_t, I_{t-1}) := \{L \in S^q_\lambda(I_t) : L \cap R \in \cup_i S^i_{\lambda \cap R}(I_{t-1})\}$; that is, L is I_t-avoiding in λ and I_{t-1}-avoiding in $\lambda \cap R$. Also define $S^q_{\lambda,R}(J_t, J_{t-1})$ analogously. Note that $S^q_{\lambda,\emptyset}(I_t, I_{t-1}) = S^q_\lambda(I_t)$ but $S^q_{\lambda,\bar{R}}(I_t, I_{t-1}) = S^q_\lambda(I_{t-1})$ if $\lambda \in \bar{R}$, and similarly for J_t.

Then proposition 5.1 is a special case of

Proposition 5.2 $|S^q_{\lambda,R}(I_t, I_{t-1})| = |S^q_{\lambda,R}(J_t, J_{t-1})|$ *for all shaded diagrams* (λ, R), $t > 0$, $q \geq 0$.

We will prove proposition 5.2 by induction, using the following recursive characterisations of $|S^q_{\lambda,R}(I_t, I_{t-1}))$ and $|S^q_{\lambda,R}(J_t, J_{t-1}))$:

Lemma 5.3 *Let* (λ, R) *be a shaded Young diagram,* $p \in (\lambda \backslash R)$ *an outer corner adjacent to* R, $q \geq 1$, $t \geq 2$, $R' = R'(R \cup \{p\})$.*Then*$|S^q_{\lambda,R}(I_t, I_{t-1})| = |S^q_{\lambda\backslash\{p\},R}(I_t, I_{t-1})| + |S^{q-1}_{\lambda/p,R'/p}(I_t, I_{t-1})|$.

Proof. Break $S^q_{\lambda,R}(I_t, I_{t-1})$ into two parts according to whether the partial transversals include the distinguished outer corner p. Specifically, let $S := \{L \in S^q_{\lambda,R}(I_t, I_{t-1}) : p \notin L\} = S^q_{\lambda\backslash\{p\},R}(I_t, I_{t-1})$, and let $S' := \{L \in S^q_{\lambda,R}(I_t, I_{t-1}) : p \in L\}$. Then there is a bijection $\gamma : S' \to S^{q-1}_{\lambda/p,R'/p}(I_t, I_{t-1})$ given by $\gamma(L) := L/p$, which completes the proof. ∎

The main difficulty, as with the proof given in the preceding sections, rests with carrying through the induction for J_t.

Lemma 5.4 *Let* (λ, R) *be a shaded Young diagram,* $p \in (\lambda \setminus R)$ *an outer corner adjacent to* R, $q \geq 1$, $t \geq 2$, $R' = R'(R \cup \{p\})$.*Then*$|S^q_{\lambda,R}(J_t, J_{t-1})| = |S^q_{\lambda\backslash\{p\},R}(J_t, J_{t-1})| + |S^{q-1}_{\lambda/p,R'/p}(J_t, J_{t-1})|$.

Proof. For convenience, let $B := S^q_{\lambda\backslash\{p\},R}(J_t, J_{t-1})$ and $C := S^{q-1}_{\lambda/p,R'/p}(J_t, J_{t-1})$. As in the proof of lemma 5.3, let $S := \{L \in S^q_{\lambda,R}(J_t, J_{t-1}) : p \notin L\}$ and $S' := \{L \in S^q_{\lambda,R}(J_t, J_{t-1}) : p \in L\}$. Then, as before, there is a bijection $\gamma : S' \to S^{q-1}_{\lambda/p,R/p}(J_t, J_{t-1})$ given by $\gamma(L) := L/p$, and $S \subseteq B$, while $C \subseteq \gamma(S')$. To complete the proof, we therefore need to show that the overcount $\mathcal{U} := B \setminus S$ is in bijection with the undercount $\mathcal{L} := \gamma(S') \setminus C$. We will prove this as a sublemma, from which the lemma will then follow immediately. Note the contrast with lemma 5.3 in which there was no corresponding overcount or undercount. ∎

Sublemma 5.5 *There exists a bijection* $g : \mathcal{U} \to \mathcal{L}$ *where* $\mathcal{U} := S^q_{\lambda\backslash\{p\},R}(J_t, J_{t-1}) \setminus \{L \in S^q_{\lambda,R}(J_t, J_{t-1}) : p \notin L\}$ *and* $\mathcal{L} := S^{q-1}_{\lambda\backslash p,R\backslash p}(J_t, J_{t-1}) \setminus S^{q-1}_{\lambda/p,R'/p}(J_t, J_{t-1})$.

6 The Second Proof (the Bijection)

Thus we have reduced the proof of the main lemma to establishing the isomorphism g of sublemma 5.5. This will be done in four steps. First we show that in any $U \in \mathcal{U}$, in a strong sense there is a unique minimal J_t appearance $l(U)$ in U; and for any $L \in \mathcal{L}$, there is a unique maximal J_{t-1} appearance $u(L) \subseteq R'$. Then we show how to 'interchange' $l(U)$ and $u(L)$ in such a way that we get a new PT $g(U) \subseteq \lambda \,/\, p$ from U and a new PT $h(L) \subseteq \lambda \setminus \{p\}$ from L. Next we prove that in fact $g(U) \in \mathcal{L}$ and $h(L) \in \mathcal{U}$. Finally, we note that g and h may be regarded as inverses to each other.

We may assume that $\mathcal{U} \neq \emptyset$ or $\mathcal{L} \neq \emptyset$, whence $p = (a, b)$, say, with $a \geq t-1$ and $b \geq t-1$. Moreover, let $R = R'(u, v)$, where without loss of generality p lies below R, i.e., $u < b$ and $v \geq a$.

Fix $U \in \mathcal{U}$. Let $A = A(U) := \{$appearances of J_n in U (with respect to $\lambda \setminus \{p\})\}$. Since $U \notin S$, $A \neq \emptyset$. Write each $\alpha \in A$ as $\{(\alpha_1 \dots \alpha_t) = (a_{\alpha 1}, b_{\alpha 1}) \dots (a_{\alpha n}, b_{\alpha t})\}$. For $i = 1 \dots t$, let $A_i := \{\alpha_i : \alpha \in A\}$.

Since $U \in S^q_{\lambda\setminus\{p\},R}(J_t, J_{t-1})$, we have $b_{\alpha 1} = b$ and $a_{\alpha t} = a$. Thus all α_1 are equal, and so are all α_t. Put $l(U)_1 = (a_1, b_1) := \alpha_1$ and $l(U)_t = (a_t, b_t) := \alpha_t$ for some $\alpha \in A$.

If some A_i were not a chain, then there were $\alpha, \beta \in A$ such that $\alpha_i <' \beta_i$. Then $(\alpha_1 \dots \alpha_i, \beta_i \dots \beta_{t-1})$ were an appearance of J_t contained in $R'(a_{\beta,t-1}, b_1)$, where $a_{\beta,t-1} < a$ but $b_1 = b$, contradicting $U \in S^q_{\lambda\setminus\{p\},R}(J_t, J_{t-1})$. Thus each A_i is a chain, and thus has a smallest element. Put $l(U)_i = (a_i, b_i) := \min A_i$, and $l(U) := (l(U)_1 \dots l(U)_t)$. Now $l(U)$ is an antichain, and $l(U)_1 <' l(U)_2 <' \dots <' l(U)_t$; for, if e.g. $l(U)_i = \alpha_i$ and $l(U)_{i+1} = \beta_{i+1}$, then $a_i \leq a_{\beta,i} < a_{\beta,i+1} = a_{i+1}$, and $b_i = b_{\alpha,i} > b_{\alpha,i+1} \geq b_{i+1}$. Thus indeed $l(U) \subset R'$ is an appearance of J_t.

Similarly, fix $L \in \mathcal{L}$; let $B := \{$appearances of J_{t-1} in $L \cap R'(a-1, b-1)$ (with respect to $(\lambda \,/\, p, R'))\}$, and let $B_i := \{i$'th components of elements in $B\}$. Then each B_i is a non-empty and finite ordered set. Put $u(L)_i = (c_i, d_i) := \max B_i$, and $u(L) := (u(L)_1 \dots u(L)_{t-1})$. Then $u(L) \subseteq R'(a-1, b-1)$ is an appearance of J_{t-1}.

With the notations as above, let $ll(U) = (ll(U)_1 \dots ll(U)_{t-1}) = ((a_1, b_2), (a_2, b_3) \dots (a_{t-1}, b_t))$ and $uu(L) = (uu(L)_1 \dots uu(L)_t) = ((c_1, b), (c_2, d_1) \dots (c_{t-1}, d_{t-2}), (a, d_{t-1}))$.

Put $U' = U \setminus l(U)$, $U'' = U' \cup ll(U)$, and $g(U) = U'' \,/\, p$. Let L' be the expansion of L, $L'' = L' \setminus u(L)$, and $h(L) = L'' \cup uu(L)$. It remains to prove that $g(U) \in \mathcal{L}$ and $h(L) \in \mathcal{U}$, and that indeed $uu(U) = l(g(U))$, $ll(L) = u(h(L))$, $h(g(U)) = U$, and $g(h(L)) = L$. Clearly, the contraction and the expansion cause no problems; whence we consider U'' rather than $g(U)$, and L' rather than L.

We start with a couple of auxiliary claims. For smoothness of the formulation and proof of the first one we define 'formal limitation points' $ll(U)_0 = (-1, b)$ and $ll(U)_t = (a, -1)$ of the anti-chain $ll(U)$, and $a_0 = b_{t+1} = -1$.

Claim 1. If $0 \leq i < j \leq t$, and $\mu_1 = (e_1, f_1) \dots \mu_t = (e_t, f_t) \in U''$ fulfil $ll(U)_i <' \mu_1 <' \mu_2 <' \dots <' \mu_t <' ll(U)_j$, then $t \leq j - i - 1$; and if $z = j - i - 1$, then $\mu_l \leq ll(U)_{i+l}$ for $l = 1 \dots z$.

Proof. We may make induction with respect to $j - i$. Thus in particular we may assume $\{\mu_1 \dots \mu_z\} \cap ll(U) = \emptyset$; i.e., $\mu_1 \dots \mu_z \in U$.

$z < t$, since $U \in S^q_{\lambda\setminus\{p\}}(J_t)$ and $R(\mu_1, \mu_z) \subseteq \lambda \setminus \{p\}$. Thus, if $z \geq j - i$, then $0 < i$ or $j < t$; and if e.g. the former, then $\{l(U)_2 \dots l(U)_i, \mu_1 \dots \mu_z, l(U)_j \dots l(U)_{t-z+j-i}\}$ were an appearance of J_t in $U \cap R'((a, b-1))$, contradiction. Thus indeed $z \leq j - i - 1$.

Now assume $z = j - i - 1$, but $\mu_\ell \not< ll(U)_{i+\ell}$ for some ℓ. If μ_ℓ and $ll(U)_{i+\ell}$ were incomparable (w.r.t. $<$), then either $\mu_\ell <' ll(U)_{i+\ell}$ or $ll(U)_{i+\ell} <' \mu_\ell$, in either case contradicting the hypothesis of induction (since $j - i - \ell < j - i$ and $i + \ell - i = \ell < j - i$). Thus in fact $ll(U)_{i+\ell} < \mu_\ell$ for each ℓ such that $\mu_\ell \not< ll(U)_{i+\ell}$.

Let $r = \min_{ll(U)_{i+\ell} < \mu_\ell} \ell$, and $s = \max_{ll(U)_{i+\ell} < \mu_\ell} \ell$. Then either $r = 1$, whence $ll(U)_i <' \mu_r \implies b_{i+r} = b_{i+1} > f_r$; or $r > 1$, whence $ll(U)_{i+r-1} > \mu_{r-1} <' \mu_r \implies b_{i+r} > f_{r-1} > f_r$; and in either case $ll(U)_{i+r} < \mu_r \implies a_{i+r} < e_r$. Thus $l(U)_{i+r} <' \mu_r$. Similarly, $\mu_s <' l(U)_{i+s+1}$. Hence, $(l(U)_2 \dots l(U)_{i+r}, \mu_r \dots \mu_s, l(U)_{i+s+1} \dots l(U)_n$ were an appearance of J_n in $U \cap R'(a, b-1)$, contradicting $U \in S^q_{\lambda\setminus\{p\}}(J_t)$.

Thus indeed $z = j - i - 1 \implies (\mu_\ell \le ll(U)_{i+\ell}$ for $\ell = 1 \dots z)$. ∎

Claim 2. If $\hat{p} \in \lambda \setminus \{p\}$ with $\hat{p} \not< p$, and $U \cap R'(\hat{p}) \in S^{|U \cap R'(\hat{p})|}_{R'(\hat{p})}(J_t)$ for some t, then $U'' \cap R'(\hat{p}) \in S^{|U'' \cap R'(\hat{p})|}_{R'(\hat{p})}(J_t)$.

Proof. Assume to the converse that there were an anti-chain $\mu_1 \dots \mu_z \in U'' \cap R'(\hat{p})$, with $\mu_\ell = (e_\ell, f_\ell)$. Then some $\mu_\ell \in U'' \setminus U = ll(U)$. Let $r = \min_{\ell \in ll(U)} \ell$, and $s = \max_{\ell \in ll(U)} \ell$. Moreover, $\hat{p} \neq p \not< \hat{p} \implies (p <' \hat{p}$ or $\hat{p} <' p)$; assume the former. Then $\mu_r = ll(U)_i$, say.

If $e_z < a$, then (by claim 1 for the sequence $\mu_{r+1} \dots \mu_z$) $z - r \le t - i - 1$; whence $(\mu_1 \dots \mu_{r-1}, l(U)_i \dots l(U)_{z+i-r})$ were a J_z appearance in $U \cap R'(\hat{p})$, contradiction. Thus $e_z > a$, whence there were ℓ in $\{s+1 \dots z\}$ such that $\mu_\ell \not\le ll(U)_\nu$ for all ν. Let v be the smallest such ℓ. Say $b_j > f_v > b_{j+1}$. Then there were a ν such that $\mu_{v-1} \le ll(U)_\nu = (a_\nu, b_{\nu+1})$, but $f_{v-1} > f_v > b_{j+1}$; whence $\nu \le j-1$; whence $e_{v-1} \le a_{j-1} < a_j$; whence $\mu_{v-1} <' ll(U)_j$. Thus (by claim 1 for the sequence $\mu_{r+1} \dots \mu_{v-1}$)

$$v - 1 - r \le j - i - 1\,. \tag{1}$$

Moreover, $\mu_v \not\le ll(U)_{j-1} \implies e_v > a_{j-1} \implies ll(U)_{j-1} <' \mu_v$; whence $\mu_v \not<' ll(U)_j$ (by claim 1 for any anti-chain in U'' between $ll(U)_{j-1}$ and $ll(U)_j$); whence $e_v > a_j$; whence $l(U)_j <' \mu_v$. Thus and by (1),

$$(\mu_1 \dots \mu_{r-1}, l(U)_{i+1} \dots l(U)_{v+i-r}, \mu_v \dots \mu_z)$$

were a J_z appearance in $U \cap R'(\hat{p})$, contradicting the assumptions. ∎

Similarly, one proves the corresponding claims 3 and 4 below.

Let $uu(L)_0 := (-1, b+1)$ and $uu(L)_{n+1} := (a+1, -1)$. Then we have

Claim 3. If $0 \le i < j \le t+1$, and $\mu_1 = (e_1, f_1) \dots \mu_z = (e_z, f_z) \in h(L)$ fulfil $uu(L)_i <' \mu_1 <' \mu_2 <' \dots <' \mu_t <' uu(L)_j$, then $z \le j - i - 1$; and if $z = j - i - 1$, then $\mu_l \ge uu(L)_{i+l}$ for $l = 1 \dots t$. ∎

Claim 4. If $\hat{p} \in \lambda \setminus \{p\}$ with $\hat{p} \not< p$, and $L' \cap R'(\hat{p}) \in S^{|L' \cap R'(\hat{p})|}_{R'(\hat{p})}(J_t)$ for some z, then $h(L) \cap R'(\hat{p}) \in S^{|h(L) \cap R'(\hat{p})|}_{R'(\hat{p})}(J_z)$. ∎

Thus, if $e_v < a_{j+1}$, then $(l(U)_2 \dots l(U)_j, \mu_v, l(U)_{j+1} \dots l(U)_n)$ were a J_t appearance, contradicting $U \in S^q_{R'((a,b-1))}(J_t)$.

Thus in fact $e_v > a_{j+1}$, whence $l(U)_{j+1} <' \mu_v$.

Next, we prove that $g(U) \in \mathcal{L}$. Let g' be the minimal pre-image of $g(U)$ under $*/p$; clearly $g' = (U \setminus l(U)) \cup ll(U)$, and first we should prove that $g' \in S^{q-1}_{\lambda,R}(J_t, J_{t-1})$.

Begin by noting that there is a 'low' and a 'high' part of $g'' := U \setminus l(U)$: Put $(\mathbf{N} \times \mathbf{N})' := \{p = (x,y) \in \mathbf{N} \times \mathbf{N} : p \le ll(U)_i$ for some $i\}$, and $(\mathbf{N} \times \mathbf{N})'' := \{p = (x,y) \in \mathbf{N} \times \mathbf{N} : p \ge l(U)_i$ for some $i\}$. Then $(\mathbf{N} \times \mathbf{N})' \cap (\mathbf{N} \times \mathbf{N})'' = \emptyset$, and g'' is the union of its low part $g'' \cap (\mathbf{N} \times \mathbf{N})'$ and its high part $g'' \cap (\mathbf{N} \times \mathbf{N})''$.

Now assume that $\mu = (\mu_1 \dots \mu_{t-1}) = ((e_1, f_1) \dots (e_{t-1}, f_{t-1}))$ (with $\mu_1 <' \mu_2 <' \dots <' \mu_{t-1}$) were a J_{t-1} appearance in g', with $R(\mu) \subseteq \lambda \cap R$. Since $g'' \in S^{q-t}_{\lambda,R}(J_t, J_{t-1})$, we have $\mu \cap ll(U) \neq \emptyset$. Say $\mu \cap ll(U) = \{\nu_1 \dots \nu_m\}$, where $\nu_1 <' \dots <' \nu_m$, and $\nu_\ell = \mu_{i_\ell} = ll(U)_{j_\ell} = (a_{j_\ell+1}, b_{j_\ell})$ for $\ell = 1 \dots m$. [[Let $\mu_i = ll(U)_j = (a_{j+1}, b_j)$ be the minimal element in $\mu \cap ll(U)$ w. r. t. $<'$.]] Moreover, let $r := \max(i : \mu_i \in (\mathbf{N} \times \mathbf{N})') \ge i_m$.

$(l(U)_1 \dots l(U)_{j_m}, \mu_{i_m+1} \dots \mu_r)$ is an anti-chain in $\lambda \cap R(a, b-1)$ and thus of length $\le t-1$, whence $r - i_m + j_m + 1 \le t$. [[Now $i_1 < j_1$, since else $(\mu_1 \dots \mu_{i_1-1}, l(U)_{i_1+1} \dots l(U)_t) \subset R$ were a J_{t-1} appearance in U, contradiction.]] If $i_{\ell+1} - i_\ell > j_{\ell+1} - j_\ell$, then

$$(l(U)_1 \dots l(U)_{j_\ell}, \mu_{i_\ell+1} \dots \mu_{i_{\ell+1}-1}, l(U)_{j_{\ell+1}+1} \dots l(U)_t)$$

were an anti-chain of length $\ge t$ in $U \cap R'(a,b)$; whence by the minimality of $l(U)$ in A we would have $\mu_{i_\ell+1} \dots \mu_{i_{\ell+1}-1} \in (\mathbf{N} \times \mathbf{N})''$, whence we could insert e.g. $l(U)_{j_\ell+1}$ into the anti-chain, but remove $l(U)_1$, getting a forbidden J_t appearance in U. [[Thus $i_m < j_m$, whence $e_{m_i} \ge a_{m_i+2}$.]] Similarly, if $f_r > b_{r-i_m+j_m+1}$, then

$$(l(U)_1 \dots l(U)_{j_m}, \mu_{i_m+1} \dots \mu_r, l(U)_{r-i_m+j_m+1} \dots l(U)_t) \in A,$$

contradiction. Thus in fact $i_m - i_1 \le j_m - j_1$ and $f_r \le b_{r-i_m+j_m+1} \le b_{r-i_1+j_1+1}$, whence

$$(\mu_1 \dots \mu_{i_1-1}, l(U)_{j_1+1} \dots l(U)_{r-i_1+j_1+1}, \mu_{r+1} \dots \mu_{t-1})$$

is a J_{t-1} appearance in $U \cap \lambda \cap R$, contradiction.

Thus $g' \cap R$ avoids J_{t-1} in R'.

Similarly, if $\mu = (\mu_1 \dots \mu_t) = ((e_1, f_1) \dots (e_t, f_t))$ is a J_t appearance in g' (w.r.t. λ), then $p \notin R(\mu)$, whence $f_1 \le b-1$ or $e_t \le a-1$. In the first case, we may proceed as before: $(\nu_1 \dots \nu_m) := \mu \cap ll(U) \ne \emptyset$ with $\nu_\ell = \mu_{i_\ell} = ll(U)_{j_\ell}$; $r := \max(i : \mu_i \in (\mathbf{N} \times \mathbf{N})')$; $r - i_1 + j_1 + 1 \le t$; $i_{\ell+1} - i_\ell \le j_{\ell+1} - j_\ell$; and $f_r \le b_{r-i_m+j_m+1}$; whence

$$(\mu_1 \dots \mu_{i_1-1}, l(U)_{j_1+1} \dots l(U)_{r-i_1+j_1+1}, \mu_{r+1} \dots \mu_t)$$

would contradict the assumption that U avoids J_t in λ.

The case $e_t \le a-1$ is excluded by symmetrical arguments.

Thus $g' \in S^{q-1}_{\lambda,R'}(J_t, J_{t-1})$, whence $g(U) \in S^{q-1}_{\lambda\,/\,p}(J_t, J_{t-1})$. On the other hand, $ll(U) \subseteq g(U) \cap R'$ shows that $g(U) \notin S^{q-1}_{\lambda\,/\,p, R'\,/\,p}(J_t, J_{t-1})$. Thus $g(U) \in \mathcal{L}$, indeed.

Now we return to the establishment of the isomorphism.

First, note that clearly $g(U) \in \mathcal{L}$ iff $U'' \in S^{q-1}_{\lambda,R}(J_t, J_{t-1}) \setminus S^{q-1}_{\lambda,R'}(J_t, J_{t-1})$.

Let $\mu_1 \dots\ \mu_t$ be any appearance of J_t in U'', and let $\hat{p} \in \mathbf{N} \times \mathbf{N}$ be the minimal point such that $\mu_1 \dots \mu_t \in R'(\hat{p})$. Then $\hat{p} \not\le p$ by claim 1 (with $i = 0$ and $j = t$). Thus, if $\hat{p} \in \lambda$, then claim 2 would yield the existence of an appearance $\lambda_1 \dots\ \lambda_n$ of J_t in U, with $R(\{\lambda_1 \dots \lambda_t\}) \subseteq R'(\hat{p}) \subseteq \lambda$, contradicting $U \in S^{\hat{p}}_{\lambda}(J_t)$. Thus in fact $\hat{p} \notin \lambda$ for each such appearance in U''; i.e., $U'' \in S^{q-1}_{\lambda}(J_t)$.

Similarly, any appearance of J_{t-1} in $U'' \cap \lambda \cap R$ would be contained in $R'(\hat{p})$ for some $\hat{p} \in (\lambda \cap R) \setminus R'(p)$, and thus yield a similar appearance in U by claim 2; contradiction. Thus indeed $U'' \in S^{q-1}_{\lambda,R}(J_t, J_{t-1})$.

On the other hand, clearly $ll(U) \subseteq U'' \cap \lambda \cap R'$ is a J_{n-1} appearance, whence $U'' \notin S^{q-1}_{\lambda,R'}(J_t, J_{t-1})$.

Thus indeed $g(U) \in \mathcal{L}$. Moreover, if $\mu_1 \dots \mu_{t-1} \in R'\,/\,p$ is an appearance of J_{t-1} in $g(U)$, then $\mu_1 \dots \mu_{t-1} \in U'' \cap R'(p)$, whence by claim 1 $\mu_\nu \le ll(U)_\nu$ for $\nu = 1 \dots n-1$. Thus indeed $u\,(g(U)) = ll(U)$.

Finally, we consider $h(L)$. If $\mu_1 \dots \mu_t \in R'(\hat{p})$ were an appearance of J_t in $h(L)$ with $q \in \lambda \backslash \{p\}$, then we would get a contradiction to claim 3 (claim 4), if $\hat{p} < p$ (else, respectively). Similarly, an appearance of J_{t-1} in $h(L) \cap R$ would contradict claim 3. The rest is easy.

7 Acknowledgements

Guoce Xin would like to acknowledge the helpful guidance of his supervisors, Bill Chen in Tianjin and Ira Gessel at Brandeis. Julian West would like to thank Bill Chen and Svante Linusson for bringing this problem to the attention of his coauthors. West's research is supported in part by NSERC operating grant OGP0105492.

References

[1] E. Babson and J. West. The permutations $123p_4 \dots p_t$ and $321p_4 \dots p_t$ are Wilf-equivalent. *Graphs and Combinatorics*, to appear.

[2] I.G. Macdonald. *Schubert Polynomials.* Publications de LACIM 6, Université du Québec à Montréal, 1991.

[3] R. Simion and F.W. Schmidt. Restricted permutations. *Eur. J. of Combinatorics*, 6 (1985), 383–406.

[4] Z. Stankova. Classification of forbidden subsequences of length 4. *Eur. J. of Comb.*, 17 (1996), 501–517.

[5] Z. Stankova. Forbidden subsequences. *Disc. Math.* 132 (1994) 291–316.

[6] J. West and Z. Stankova. A New Class of Wilf-equivalent permutations. *J. Alg. Combin.*, to appear.

[7] J. West. *Permutations with forbidden subsequences and stack-sortable permutations.* PhD thesis, M.I.T., 1990.

[8] J. West. Generating trees and the Catalan and Schröder numbers. *Disc. Math.* 146 (1995) 247–262.

[9] J. West. Generating trees and forbidden subsequences. *Disc. Math.* 157 (1996) 363–374.