ON SOME PROPERTIES OF PERMUTATION TABLEAUX

ALEXANDER BURSTEIN

ABSTRACT. We consider the relation between various permutation statistics and properties of permutation tableaux. We answer some of the questions of Steingrímsson and Williams [9], in particular, on the distribution of the bistatistic of numbers of rows and essential ones in permutation tableaux. We also consider and enumerate sets of permutation tableaux related to some pattern restrictions on permutations.

1. INTRODUCTION

Permutation tableaux are a particular class of J-diagrams that were studied by Postnikov [7] and enumerated by Williams [10]. They are defined as follows. Given a partition of $\lambda = (\lambda_1 \ge \lambda_2 \ge \cdots \ge \lambda_k > 0)$ of an integer $m = \sum_i \lambda_i$, its Young diagram Y_{λ} of shape λ is a left-justified diagram of m boxes, with λ_i boxes in row i.

A permutation tableau \mathcal{T}_n^k , which we will also call a 1-hinge tableau, is a partition λ whose Young diagram Y_{λ} is contained in a $k \times (n-k)$ rectangle aligned with its top and left edges, together with a filling of the cells of Y_{λ} with 0s and 1s that satisfies the following properties:

(column): Each column of the rectangle contains at least one 1.

(1-hinge): A cell in Y_{λ} with a 1 above it in the same column and a 1 left of it in the same row must contain a 1.

The filling satisfying the column and hinge properties is called *valid*. The 1s in \mathcal{T}_n^k that are topmost in their columns or leftmost in their rows are called *essential*, and the remaining 1s, i.e. those that are forced by the 1-hinge property, are called *induced*. Note that the column property implies that Y_{λ} must have exactly n-k columns and at least k rows. Note also that some rows of \mathcal{T}_n^k (as opposed to columns) may contain all zeros. Removing the hinge requirement yields the definition of J-diagram [7].

Alternatively, a permutation tableau \mathcal{T}_n^k may be thought of as a filling of $k \times (n-k)$ rectangle with 0s, 1s and 2s such that the cells inside Y_λ are filled with 0s and 1s so as to satisfy the column and 1-hinge properties, and the cells outside Y_λ are filled with 2s.

Properties of permutation tableaux were studied by Steingrímsson and Williams [9]. They gave a simpler description of a map described by Postnikov [8] that takes permutation tableaux contained in a $k \times (n-k)$ rectangle to permutations in S_n with k weak excedances, and proved that this map Φ is a bijection that also preserves many other statistics. The bijection Φ will be described in the next section.

One of the conjectures made in [9, Section 7] is that the distribution of permutation tableaux according to the number of essential 1s is equal to that for number of cycles in permutations, i.e. it is given by the signless Stirling numbers of the first kind. Moreover, [9] conjectured that the joint distribution of tableaux according to the number of rows and the number of essential 1s equals that of permutations according to the number of weak excedances and the number of cycles of a permutation, when written in standard cycle form (which is the same distribution as that of permutations according to descents and left-to-right-minima). In this paper we will give a simple natural bijection on J-tableaux that induces the conjectured bijection above and preserves several other statistics.

Another bijection Ψ on \mathfrak{S}_n defined in [9] translates certain statistics on permutations corresponding to entries of their permutation tableaux (determined by Φ) into certain linear combinations of generalized

Date: May 8, 2006.

Research supported in part by the NSA Young Investigator Grant H98230-06-1-0037.

ALEXANDER BURSTEIN

permutation patterns. In particular, Ψ maps permutations avoiding generalized pattern 2-31 to permutation tableaux with exactly one 1 in each column. Since the number of permutations in \mathfrak{S}_n avoiding 2-31 is C_n , the *n*th Catalan number, we give a simple description of permutations whose tableau has a single 1 in each column in terms of noncrossing partitions.

Finally, we describe the properties of the tableaux of permutations restricted by some 3-letter patterns and enumerate those of them that have the maximal number of essential 1s.

2. Permutations and permutation tableaux

Here we briefly describe a bijection Φ from [9]. We also give a different proof that Φ is a bijection, essentially recovering the tableaux from the corresponding permutation by reconstructing the columns of the tableau from left to right (or, similarly, rows from top to bottom), as opposed to right-to-left column construction in [9] (or a similar bottom-to-top row construction).

Given a permutation tableau \mathcal{T}_n^k , its diagram $D(\mathcal{T}_n^k)$ is defined as follows. The southeast border of partition λ gives a path $P = (P_i)_{i=1}^n$ of length n from the northeast corner of the $k \times (n-k)$ rectangle containing λ to the southwestern corner of that rectangle. Label each step P_i in P with i, for $i \in [n]$ (where $[n] = \{1, 2, \ldots, n\}$). Now, given an edge P_i , also label with i the edge Q_i on the opposite end of the row (if P_i is vertical) or column (if P_i is horizontal) containing the edge P_i . Replace each 1 in \mathcal{T}_n^k with a vertex and delete all 0s. From each vertex draw edges east and south either to the closest vertices in the same row and the same column or to the labels i of some edge P_i in P. The resulting picture is the diagram $D(\mathcal{T}_n^k)$. It is also convenient to consider $D(\mathcal{T}_n^k)$ together with the edges from the labels i of edges Q_i on the northwestern boundary of \mathcal{T}_n^k to the closest (leftmost) vertex in the same row (if Q_i is vertical) or the closest (rightmost) vertex in the same column (if Q_i is horizontal). We will denote the resulting diagram $D'(\mathcal{T}_n^k)$ and call it the expanded diagram of \mathcal{T}_n^k . It is also convenient to think of edge labels $i \in [n]$ as labeling a row or a column between P_i and Q_i , and label each cell in the tableau by the ordered pair of its row and column labels.

Given a tableau \mathcal{T}_n^k as above, the permutation $\pi = \Phi(\mathcal{T}_n^k)$ as defined as follows. For each $i \in \{1, \ldots, n\}$, consider a zigzag path in $D'(\mathcal{T}_n^k)$ that starts at Q_i (going south or east depending on whether Q_i is horizontal or vertical) and switchings direction between south and east at every vertex it encounters. If that path terminates at P_j , then we set $\pi(i) = j$. Alternatively, in $D(\mathcal{T}_n^k)$ we replace the first edge of the path starting from Q_i by an edge north or west from P_i to the vertex in row or column *i* that is farthest from P_i (and closest to Q_i).

It is shown in [9] that Φ is a bijection, $\pi(i) \geq i$ if P_i is vertical and $\pi(i) < i$ if P_i is horizontal, hence $\pi = \Phi(\mathcal{T}_n^k)$ has k weak excedances (positions i such that $\pi(i) \geq i$) and n - k deficiencies (positions i such that $\pi(i) < i$). In fact, i is a fixed point of π (i.e. $\pi(i) = i$) if and only if the row of \mathcal{T}_n^k labeled i does not contain a 1. We will refer to Φ as the standard bijection from permutations tableaux to permutation and refer to $\Phi(\mathcal{T})$ as the permutation of \mathcal{T} and to $\Phi^{-1}(\pi)$ as the tableau of π .

In other words, π is a derangement if and only if every row of $\mathcal{T} = \Phi^{-1}(\pi)$ also must contain a 1. Therefore, in this we may reflect the tableau \mathcal{T} and obtain another tableau $\operatorname{ref}(\mathcal{T})$ with the same properties. Moreover, the each southeast path $Q_i(\mathcal{T}) \to P_{\pi(i)}(\mathcal{T})$ is thus reflected onto a southeast path $Q_{n+1-i}(\operatorname{ref}(\mathcal{T})) \to$ $P_{n+1-\pi(i)}(\operatorname{ref}(\mathcal{T}))$. Thus, if $\pi = \Phi(\mathcal{T})$ and $\sigma = \Phi(\operatorname{ref}(\mathcal{T}))$, then $\sigma(n+1-i) = n+1-\pi(i)$ for all $i \in [n]$, i.e. σ is the reversal of the complement of π . This answers Open Problem 5 of [9].

Before proceeding further we will introduce a bit of terminology.

Definition 2.1. If $\pi(i) \ge i$ is a weak excedance, we will call $\pi(i)$ a weak excedance top (wex top) of π , and call *i* a weak excedance bottom (wex bottom) of π . If $\pi(i) < i$, we will call $\pi(i)$ a non-weak-excedance bottom (nonwex bottom) of π , and call *i* a non-weak-excedance top (nonwex top) of π .

We will now describe a way to recover \mathcal{T}_n^k from $\pi = \Phi(\mathcal{T}_n^k) \in \mathfrak{S}_n$ starting from the leftmost column. Let m be the largest position of a non-fixed point in π . Then $\pi(m) < m$, and $\pi(i) = i$ for all i > m. Since it is trivial to recover the edges P_i for i > m, we may assume that without loss of generality that m = n (so P_n is a horizontal edge).

Theorem 2.2. Assume that $\pi \in \mathfrak{S}_n$ is such that $\pi(n) < n$, and let $\mathcal{T} = \Phi^{-1}(\pi)$. Suppose that the dots in the leftmost column of \mathcal{T} are in rows labeled $i_1 < i_2 < \cdots < i_r$. Furthermore, let \mathcal{T}' be the tableau obtained by removing the leftmost column of \mathcal{T} and replacing P_n with a vertical edge (so that n becomes a fixed point), and let $\pi' = \Phi(\mathcal{T}')$. Then $\pi = \pi' \circ (i_1 i_2 \dots i_r n)$, $\pi(n) < \pi(i_1) < \dots \pi(i_r)$ and $\pi(i_1) < \dots \pi(i_r)$ are the successive non-fixed-point left-to-right maxima of the subsequence of π consisting of values greater than $\pi(n)$.

Proof. For $1 \leq j < r$, the first three steps of the southeast path from Q_{i_j} to $P_{\pi(i_j)}$ are east from Q_{i_j} to cell (i_j, n) , then south to (i_{j+1}, n) , then east from (i_{j+1}, n) . Thus, $\pi(i_j) = \pi'(i_{j+1})$. Similarly, the path starting from Q_{i_r} goes east to (i_r, n) , then south to P_n , so $\pi(i_r) = n = \pi'(n)$. Likewise, the path from Q_n starts south to (i_1, n) , then turns east, so $\pi(n) = \pi'(i_1)$. Thus, $\pi = \pi' \circ (i_1 i_2 \ldots i_r n)$ as claimed.

Note that paths $p(i_j): i_j \to \pi(i_j)$ and $p(i_{j+1}): i_{j+1} \to \pi(i_{j+1})$ meet at a vertex in cell (i_{j+1}, n) , where $p(i_j)$ enters it traveling south and leaves east while $p(i_{j+1})$ enters it traveling east and leaves south. Hence, we can see by induction that at each row $p(i_{j+1})$ is to the east of $p(i_j)$ and at each column $p(i_j)$ is to the south of $p(i_{j+1})$, so if $p(i_j)$ and $p(i_{j+1})$ meet again at a cell it must contain an induced 1, so they cannot cross. It follows from that $\pi(i_j) < \pi(i_{j+1})$ for all j < r, and similarly that $\pi(n) < \pi(i_1)$.

Now let $l \in [n]$ be a row label such that $i_{j-1} < l < i_j$ (if j = 1, we simply let $l < i_1$). Then the path from Q_l starts east and either continues to P_l or first turns south at a cell (l, m) for some m < n. Since (l, m) (if it exists) is northeast of (i_{j+1}, n) , the same argument as before applies to show that $\pi(i_j) > \pi(l)$. Therefore, $\pi(i_{j+1})$ is the leftmost value to the right of $\pi(i_j)$ that is greater than $\pi(i_j)$, so the theorem follows.

Note that largest entry n in the (increasing) cycle above gives the label of the leftmost column, while the remaining entries give the labels of the rows containing dots in that column. Iterating the operation yields all the cells in \mathcal{T} containing dots. The largest elements in cycles are column labels, the rest are row labels.

Example 2.3. Let $\pi = 65187243$. Then we have

i	1	2	3	4	5	6	7	8	cycle					
π	6	5	1	8	7	2	4	3	(148)	τ	•		•	•
π'	3	5	1	6	7	2	4	8	(2457)			•	•	•
π''	3	4	1	5	6	2	7	8	(12456)	I =	•	•	•	
$\pi^{\prime\prime\prime}$	2	3	1	4	5	6	7	8	(123)			•	•	
e	1	2	3	4	5	6	7	8				-	-	J

so $\pi = (123)(12456)(2457)(148)$, and the path $P(\pi) = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ v & v & h & v & v & h & h \end{pmatrix}$, where v and h denote the vertical and horizontal edges, respectively, so $\mathcal{T} = \Phi^{-1}(\pi)$ has dots in cells labeled (1,3), (2,3), (1,6), (2,6), (4,6), (5,6), (2,7), (4,7), (5,7), (1,8), (4,8). Note that row labels (here, 1, 2, 4, 5) increase from top to bottom, while column labels (here, 3, 6, 7, 8) increase from right to left.

Remark 2.4. Note that the column rule implies that there are no 1-cycles. Likewise, the 1-hinge rule implies that if a cycle c in the product as above contains an element i, and another cycle c' to its left contains elements j_1, j_2 such that $j_1 < i < j_2$, then c' also contains i.

Definition 2.5. We call the (unique) representation of a permutation π as a product of increasing cycles subject to the conditions in Remark 2.4 the *column decomposition* of π .

Likewise, it is easy to see that we can determine \mathcal{T} by rows from top to bottom by decomposing \mathcal{T} as a product of *decreasing* cycles. In this case, however, we may have 1-cycles, which will correspond to rows without dots. Here, at each step, we will need to find positions of the successive non-fixed-point *right-to-left minima* of the subsequence of π consisting of values smaller than $\pi(s)$, where s is the position of the leftmost non-fixed point. After arriving at the identity permutation, we add the remaining elements as fixed points.

Example 2.6. Given $\pi = 65187243$, we obtain similarly to Example 2.3 that $\pi = (765)(8764)(7632)(8631)$.

Remark 2.7. As in Remark 2.4, note that the 1-hinge rule implies that if an element i occurs in some cycle c in the product as above, and another cycle c' to the left of c contains elements j_1, j_2 such that $j_1 > i > j_2$, then c' also contains i.

ALEXANDER BURSTEIN

Definition 2.8. We call the (unique) representation of a permutation π as a product of decreasing cycles subject to the conditions in Remark 2.7 the row decomposition of π .

3. TABLEAUX OF RESTRICTED PERMUTATIONS

Now we need to define the notion of a pattern. A (classical) permutation pattern is an order-isomorphism type of a sequence of totally ordered letters. Such a sequence is then referred to as an occurrence or instance of that pattern. For example, $\sigma = 214653$ contains instances of pattern $\pi = 2$ -3-1 at subsequences 463 and 453 but no instance of pattern $\tau = 3$ -1-2. In this case, we say that σ contains π and avoids τ . The dashes, which are often dropped when referring to classical patterns, are used to indicate that the terms involved in an occurrence of the pattern may be separated by an arbitrary number of other terms. A generalized permutation pattern (introduced by [1]) is a pattern where some letters adjacent in a pattern must also be adjacent in a containing permutation. Such adjacent pairs in a pattern are then not separated by a dash. For example, 453 is an instance of pattern 2-31 in 214653, while 463 is an instance of pattern 23-1, but not 2-31. If π is a pattern and σ is a permutation, we write $(\pi)\sigma$ for the number of occurrences of π in σ . We also write $\mathfrak{S}_n(\pi) = \{\sigma \in \mathfrak{S}_n \mid (\pi)\sigma = 0\}$. One of the earliest results concerning patterns is that $|\mathfrak{S}_n(\pi)| = C_n$, the *n*th Catalan number, for any classical pattern $\pi \in \mathfrak{S}_3$. The same holds for $\pi = 2$ -31 and $\pi = 31$ -2 since $\mathfrak{S}_n(2-31) = \mathfrak{S}_n(2-3-1)$ and $\mathfrak{S}_n(31-2) = \mathfrak{S}_n(3-1-2)$. See [2] for more on patterns and pattern avoidance.

We will now show a nice application of row decomposition. In [9], another bijection $\Psi : \mathfrak{S}_n \to \mathfrak{S}_n$ is given (to be described later on in this paper) that implies that the permutation tableaux containing a single 1 in each column (i.e. the fewest possible number of 1s) are the images of permutations avoiding pattern 2-31, and thus are counted by the Catalan number $C_n = \frac{1}{n+1} \binom{2n}{n}$. Open Problem 6 in [9] asks for "a bijection from these tableaux to any well-known set of objects enumerated by Catalan numbers". We will give such a bijection using Remark 2.7.

Theorem 3.1. There is a natural bijection between k-row permutation tableaux with n - k 1s (i.e. a single 1 per column) and noncrossing partitions of [n] with k blocks.

Hence, the number of permutations corresponding to these tableaux that have k weak excedances is the number of partitions of [n] with k blocks, i.e. the Narayana number $N(n,k) = \frac{1}{n} \binom{n}{k} \binom{n}{k-1}$ for any $k \in [n]$.

Proof. Suppose that \mathcal{T} is a tableau with a single 1 in each column, and let $\pi = \Phi(\mathcal{T}) \in \mathfrak{S}_n$. Consider the row decomposition of π . Since each column contains a single 1, each column label occurs in only one of the cycles. Each row label also occurs in a single cycle. Thus, every label occurs once in the row decomposition of π , hence the cycles in the row decomposition of π are mutually disjoint, and the underlying sets for these cycles form a set partition Π of [n]. Suppose that there cycles $\gamma_1 \neq \gamma_2$ in the row decomposition of π such that γ_1 contains elements a, c and γ_2 contains elements b, d such that a > b > c > d. Thus, if γ_2 is to the left of γ_1 , then γ_2 also contains b, and if γ_1 is to the left of γ_2 , then γ_1 also contains c. This contradicts the fact that each element in [n] must occur in a single cycle of the row decomposition of π , so Π is noncrossing. \Box

We note that the set E_n of permutations whose cycle decomposition corresponds to noncrossing partitions of *n* occurs in [3] as the set of sequences of halves of even-valued entries of 3-1-4-2 avoiding Dumont permutations of the second kind. Note also that E_n is exactly the set of permutations whose column decomposition contains only 2-cycles.

The next theorem describes the tableaux of 3-2-1 avoiding permutations.

Theorem 3.2. The tableaux of 3-2-1 avoiding permutations are exactly those whose rows and columns are all nondecreasing from left to right and from top to bottom, respectively.

Proof. Note that π is a 3-2-1 avoiding permutation if and only if each element of π is either a left-to-right maximum or a right-to-left-minimum, i.e. if and only if π is the identity or a union of two nondecreasing subsequences. Again, without loss of generality assume that n is not a fixed point of π , i.e. $\pi(n) < n$. Suppose that the first column contains a 0 at row l that is underneath a 1 at row i. Then as in the proof of Theorem 2.2, we have $i < l \leq \pi(l) < \pi(i)$. Hence, to avoid an occurrence of 3-2-1 in π , we must have

 $\pi(m) > \pi(l) \ge l$ for all m > l, so π must have at least n - l + 1 values greater than $\pi(l) \ge l$ ($\pi(i)$ and all $\pi(m)$ for m > l), which is impossible. Therefore, the leftmost column of π must have all 0s atop all 1s. Moreover, if the rows containing a dot in the leftmost column are labeled i_j ($1 \le j \le r$), then the sequence $\{\pi(i_j)\}_{j=1}^r$, is increasing.

Let π' be the permutation defined as in Theorem 2.2. Assume that the sequence of wex tops of π' (see Definition 2.1) is increasing. We have $\pi(l) = \pi'(l)$ for a wex bottom $l < i_1$, as well as $\pi(i_j) = \pi'(i_{j+1})$ for $1 \leq j < r$, and $\pi(i_r) = \pi'(n) = n$, so the sequence of wex tops of π is increasing if the sequence of wex tops of π' is increasing. Since the identity tableau has no columns and increasing wex tops, we see by induction that if each column of π has all 0s atop all 1s, then the sequence of wex tops of π is increasing.

Thus, if π is 3-2-1 avoiding then the sequence of wex tops of π is increasing and each column of the tableau of π has all 0s atop of all 1s. Likewise, π is the sequence of nonwex bottoms of π is increasing and each row of the tableau of π has all 0s to the left of all 1s. Conversely, if both wex tops and nonwex bottoms are increasing, then π is a union of at most 2 subsequences and hence avoids 3-2-1.

Now we will describe a bijection Ψ on \mathfrak{S}_n that translates certain pattern statistics on permutations into alignment and crossing statistics (see [9]) of the tableaux of their images. Given a permutation $\pi \in S_n$ and we define the right embracing number REMB($\pi(i)$) of the entry $\pi(i)$ as the number of instances of 2-31 in π that start with $\pi(i)$. We also define the sequence REMB($\pi) = \{\text{REMB}(\pi(i)) | i \in [n]\}$. If $\pi(i) > \pi(i+1)$, then $\pi(i)\pi(i+1)$ is a descent of π , $\pi(i)$ a descent top and $\pi(i+1)$ is a descent bottom. Similarly, if $\pi(i) < \pi(i+1)$, then $\pi(i)$ is a non-descent bottom and $\pi(i+1)$ is a non-descent top. Note that a permutation is uniquely determined by the sets of its descents tops and descent bottoms and the sequence of its right embracing numbers [9]. For $\pi \in \mathfrak{S}_n$, define the following sets:

$DB(\pi) = \text{set of descent bottoms of } \pi,$	$NDT(\pi) = set of non-descent tops of \pi,$
$DT(\pi) = \text{set of descent tops of } \pi,$	$NDB(\pi) = set of non-descent bottoms of \pi,$
WEXB (π) = set of wex bottoms of π ,	NWEXT (π) = set of nonwex tops of π ,
WEXT (π) = set of wex tops of π ,	NWEXB (π) = set of nonwex bottoms of π .

Note that $NDT(\pi) = [n] \setminus DB(\pi)$, $NDB(\pi) = [n] \setminus DT(\pi)$, $NWEXT(\pi) = [n] \setminus WEXB(\pi)$, $NWEXB(\pi) = [n] \setminus WEXT(\pi)$. The permutation $\sigma = \Psi(\pi)$ is defined as follows. We set

> WEXB $(\sigma) = \{a+1 \mid a \in DB(\pi)\} \cup \{1\},$ NWEXT $(\sigma) = \{a+1 \mid a \in NDT(\pi), a \neq n\},$ WEXT $(\sigma) = \{a-1 \mid a \in DT(\pi)\} \cup \{n\},$ NWEXB $(\sigma) = \{a-1 \mid a \in NDB(\pi), a \neq 1\}$

Next we use $\text{REMB}(\pi)$ to determine the bijections $\text{WEXB}(\sigma) \to \text{WEXT}(\sigma)$ and $\text{NWEXT}(\sigma) \to \text{NWEXB}(\sigma)$ that together form Ψ . First, we find $a = \max \text{WEXB}(\sigma)$, then find the $(\text{REMB}_{\pi}(a) + 1)$ -st smallest element in $\text{WEXT}(\sigma)$ among those greater than or equal to a, call it b. Then $b = \sigma(a)$. Next we delete a from $\text{WEXB}(\sigma)$ and b from $\text{WEXT}(\sigma)$ and iterate this process until there are no more wex bottoms left. Similarly, we find $c = \min \text{NWEXT}(\sigma)$, then find $(\text{REMB}_{\pi}(c) + 1)$ -st largest element in $\text{NWEXB}(\sigma)$ among those less than c, call it d. Then $d = \sigma(c)$. Next we delete c from $\text{NWEXT}(\sigma)$ and d from $\text{NWEXB}(\sigma)$ and iterate this process until there are no more nonwex tops left.

Recall that [9] showed that $\Psi(\mathfrak{S}_n(2-31))$ is the set of permutations whose tableaux have a single 1 per column. Here we establish another pattern-related result concerning Ψ .

Theorem 3.3. $\Psi(\mathfrak{S}_n(31\text{-}2)) = \mathfrak{S}_n(3\text{-}2\text{-}1)$, *i.e.* $\Psi(\mathfrak{S}_n(31\text{-}2))$ is the set of permutations whose tableaux have nondecreasing rows and columns.

Proof. Let $\pi \in \mathfrak{S}_n$, and let $\sigma = \Psi(\pi)$. It is shown in [9] that the occurrences of 3-12 in π correspond to pairs $i, j \in [n]$ such that $j < i \leq \sigma(i) < \sigma(j)$ or $\sigma(j) < \sigma(i) < i < j$. In other words, i and j are both wex bottoms or both nonwex tops (i.e. $\sigma(i)$ and $\sigma(j)$ are both wex tops or both nonwex bottoms), and σ has an inversion at positions (i, j). Therefore, π avoids 31-2 if and only if the sequence of wex tops of σ and the sequence of nonwex bottoms of σ both have no inversions, i.e. are increasing. In other words, σ is a union of at most two increasing subsequences, that is $\sigma \in \mathfrak{S}_n(3-2-1)$.

ALEXANDER BURSTEIN

4. Essential 1s

Definition 4.1. Given a permutation tableau \mathcal{T} , we call the topmost 1 in each column and the leftmost 1 in each row an *essential* 1. If a 1 is both the leftmost 1 in its row and the topmost 1 in its column, then we call it *doubly essential*. Let $ess(\mathcal{T})$ and $dess(\mathcal{T})$ be the number of essential 1s and doubly essential 1s in \mathcal{T} .

Note that for any tableau \mathcal{T} in a $k \times (n-k)$ rectangle, we have $\operatorname{ess}(\mathcal{T}) + \operatorname{dess}(\mathcal{T}) = k + (n-k) = n$ by a simple sieve argument. Obviously, each nonessential 1 is induced by a pair of essential 1s so each tableau is determined by its essential 1s. It was conjectured in [9] that distribution of the bistatistic (ess, rows) of the number of essential 1s and the number of rows on permutation tableaux is the same as that of the bistatistic (cycles, wex) of the number of cycles and the number of weak excedances on permutations. We will show the following.

Theorem 4.2. The bistatistics (dess, rows) on tableaux bounded by a path of length n and (cycles, wex) on permutations in \mathfrak{S}_n have the same distribution.

Consideration of essential 1s suggests a different type of tableau, which we will call a *bare tableau* or a *0-hinge tableau*, that results from removing all nonessential 1s from a permutation tableau. A bare tableau is defined almost the same way as a permutation tableau, except that the 1-hinge property is replaced with the following 0-hinge property:

(0-hinge): A cell in Y_{λ} with a 1 above it in the same column and a 1 to its left in the same row must contain a 0.

We also define (see (4.1)) two maps θ and ϕ on tableaux of the same shape that result in removal and filling of all nonessential 1s, respectively, as well as a map Θ from bare tableaux to permutations, similar to Φ , i.e. given by southeast paths from the northwestern to the southeastern boundary of the tableau that switch direction at each 1. Empty circles denote essential 1s, double circles denote doubly essential 1s. (4.1)

$$\phi(\mathcal{B}) = \mathcal{T} = \begin{bmatrix} \textcircled{\textcircled{0}} & \textcircled{0} & \textcircled{0} \\ \hline & \rule{0} & \rule{0$$

It is easy to see that the diagram $D(\mathcal{B})$ of a bare tableau is a binary forest, since every vertex may only have a single edge connecting it to a vertex above or to the left of it and at most two edges to vertices south and east of it. Indeed, if there are edges south and east to the same vertex, then \mathcal{B} has a 1-hinge at that vertex, which is impossible. Moreover, the doubly essential 1s are exactly the roots of those binary trees.

The resulting trees are labeled as follows. The root has the two labels of the row and column of the cell that contains it. Each nonroot vertex has a single label: the column (resp. row) label of the cell containing it if that vertex is reached by an edge south (resp. east) when traveling from a root vertex. Call a nonroot vertex a *left* son if it get a row label, and a *right* son if it gets a column label. Then the following property is easy to see: both sons of each left son labeled a have labels greater than a, and both sons of each right son labeled b have labels less than b. Also, the root gets the least and the greatest label in each tree, and thus has the properties of both a left son and a right son.

Thus, each bare tableau \mathcal{B} can be decomposed into several tableaux each of which corresponds to a single binary tree in $D(\mathcal{B})$ labeled as described above.

$$(4.2) \qquad \mathcal{B} = \underbrace{\bigcirc \circ \circ}_{\circ} = \left(\bigcirc \circ \circ \circ \\ \circ & \circ \\ \circ$$

The permutation $\tau = \Theta(\mathcal{B})$ is obtained by traversing each labeled binary tree as in (4.2) as follows:

 $\Phi(\mathcal{T}) = \pi$

- (1) Start from the smallest label (at the root) along the *l*-edge, if possible. If there is no left son, this is the first return to the root (see last rule).
- (2) At each step, start at the previously reached vertex (say, v) and attempt to travel away from the root alternating unused l and r edges for as long as possible. The label of the end vertex w of this path is the next term in the cycle.
- (3) If no edges starting at v away from the root exist or are unused, travel from v towards the root along the edges with the same labels as long as possible. The label of the end vertex w of this path is the next term in the cycle.
- (4) At the first return to the root, the next term is the largest label at the root. At the second return to the root (including the case when there is no right son), the cycle is complete.

For example, the traversal of the trees in (4.2) yields the cycles (14863) and (257) as desired.

Note that the removal of nonessential 1s leaves the labels of rows and columns the same, and rows still correspond to weak excedances (in particular, zero rows correspond to fixed points), while columns correspond to deficiencies. Thus, to prove Theorem 4.2 we first need to prove the following.

Theorem 4.3. If \mathcal{B} is bare tableau with no zero rows and a single doubly essential 1, then $\tau = \Theta(B)$ is a cyclic permutation of length greater than 1. Moreover, \mathcal{B} is uniquely recoverable from τ in the cycle notation.

Proof. The doubly essential 1 (i.e. the root of the corresponding labeled binary tree) must have the labels (1, n) where n is the length of the southeast boundary P of \mathcal{B} . Suppose the successive left sons away from the are labeled i_1, i_2, \ldots, i_r . Then $1 < i_1 < i_2 < \cdots < i_r < n$, and each i_{j+1} is the left son of i_j (letting $i_0 := 1$ and $i_{r+1} := n$). Let $D(\mathcal{B}_{i_j})$ be the subtree of $D(\mathcal{B})$ with the right son of i_j as the root. Then

 $\operatorname{traversal}(D(\mathcal{B})) = 1, \operatorname{traversal}(D(\mathcal{B}_{i_1})), i_1, \operatorname{traversal}(D(\mathcal{B}_{i_2})), i_2, \dots, \operatorname{traversal}(D(\mathcal{B}_{i_r})), i_r, n, \operatorname{traversal}(D(\mathcal{B}_n))$

Each \mathcal{B}_{i_j} is strictly smaller than \mathcal{B} , so by inductive assumption the traversal of \mathcal{B}_{i_j} contains every vertex label of \mathcal{B}_{i_j} once. Hence, the traversal of \mathcal{B} contains every label of \mathcal{B} once, i.e. $\tau = \Theta(\mathcal{B})$ is a cyclic permutation.

Note also that the first part of Theorem 2.2 (i.e. that $\pi = \pi' \circ (i_1 i_2 \dots i_r n)$) remains true for bare tableaux after changing Φ to Θ . Thus, to show that Θ is a bijection we only need to prove that we can uniquely recover the leftmost column of the bare tableau \mathcal{B} with no fixed points and a single doubly essential 1 from the cyclic permutation $\tau = \Theta(\mathcal{B})$.

It is not hard to see that because i_j is a left son, every label in $D(\mathcal{B}_{i_j})$ is larger than i_j . Hence, if ρ is obtained by removing the parentheses from τ written in the cycle notation starting with 1, then $\{i_r, i_{r-1}, \ldots, i_1, 1\}$ are the right-to-left minima of the prefix of ρ before n.

Now Theorem 4.3 clearly implies Theorem 4.2. Note that we can similarly prove that the column labels of vertices in the top row (i.e. row 1) of \mathcal{B} are the right-to-left maxima of the suffix of ρ starting with n.

References

- E. Babson, E. Steingrímsson, Generalized permutation patterns and a classification of the Mahonian statistics, Sém. Lothar. Combin., B44b (2000), 18 pp.
- [2] M. Bóna. Combinatorics of Permutations, Chapman & Hall/CRC Press, 2004.
- [3] A. Burstein, S. Elizalde, T. Mansour, Restricted Dumont permutations, Dyck paths, and noncrossing partitions, to appear in *Proceedings of the 18th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC'06)*, June 2006, San Diego, CA, USA. Online at http://garsia.math.yorku.ca/fpsac06/papers/78_ps_or_pdf.pdf.
- [4] A. Claesson: Generalized pattern avoidance, Europ. J. Combin. 22 (2001), 961–971.
- [5] S. Corteel, Crossing and alignments of permutations, preprint, arXiv:math.CO/0601469, 13 pp.
- [6] S. Corteel, Crossings and alignments of permutations, preprint, arXiv:math.CO/0505031, 11 pp.
- [7] A. Postnikov, Webs in totally positive Grassmann cells, in preparation.
- [8] A. Postnikov, personal communication.
- [9] E. Steingrímsson, L.K. Williams, Permutation tableaux and permutation patterns, preprint, arXiv:math.CO/0507149.
- [10] L. Williams, Enumeration of totally positive Grassmann cells, Advances in Math., 190 (2005), 319–342.

DEPARTMENT OF MATHEMATICS, IOWA STATE UNIVERSITY, AMES, IA 50011-2064, USA

E-mail address: burstein@math.iastate.edu

URL: http://www.math.iastate.edu/burstein/