Representing permutations with tableaux

Ander Erickson

May 6, 2007

Under the direction of Professor John Caughman

A math 501 project submitted in partial fulfillment of the requirements for the degree of Masters of Science in Mathematics

The purpose of this paper is to explore and restate some of the results put forth by Steingrimsson and Williams in their article "Permutation Tableaux and Permutation Patterns" [1]. The concept of a permutation tableau, which appeared previously only in a unpublished work of Postnikov [2], is a variation of a Young diagram in which the boxes are filled only with 0's, 1's, and 2's, according to a few simple rules. Postnikov found a bijection between permutation tableaux and permutations (another bijection also appears in a paper by Williams [3]), and in Steingrimsson and Williams, a simplified bijection was presented. The verification that the simple mapping is a bijection, however, appeared unnecessarily complex. In this paper, I present the same bijection with a proof of even greater simplicity. I then continue to follow the article, showing how different characteristics (the number and type of crossings) of permutations are reflected in their equivalent tableaux. A second bijection is then defined that highlights a symmetry between certain characteristics of the diagrams and certain permutation patterns that involve descents.

1 What is a Permutation Tableau?

1.1 The Basic Definition

A permutation tableau is a *Young diagram* with 1's and 0's inserted according to the simple rules given below. What's a Young diagram? Well, begin with a finite collection of non-negative integers and put them in a weakly decreasing order. Let's call these integers $a_1 \geq a_2 \geq \cdots \geq a_k$. Now construct a horizontal row of a_1 boxes, place a horizontal row of a_2 boxes below that (make sure that the rows are aligned on the left), and continue to follow the same procedure until you've finished. Figure 1 displays a Young diagram using the sequence 5,3,3,2,1 and another with the sequence 3,3,3,0,0.

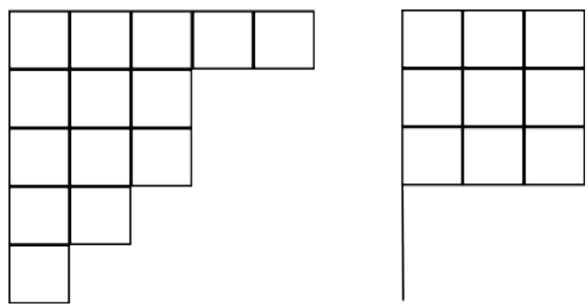

Figure 1: Examples of Young diagrams

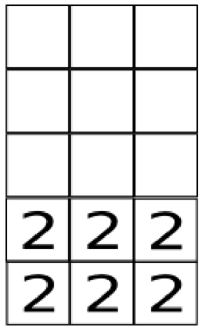

Figure 2: Young diagram with 2's inserted

If you have one or more zeroes in your set, then you will have to come up with some way of indicating the number of empty rows that are on the bottom. The second Young diagram in Figure 1 gives an indication of how hard it is to parse through the use of left-aligned edges. Another way of doing so is as follows: embed the diagram in an a_1 by k grid and fill in the boxes that aren't part of the Young diagram with 2's (these are simply serving as place-holders – we're using 2's because 0's and 1's will be used for something else shortly). Figure 2 represents the second diagram from Figure 1 with 2's inserted to clarify how many rows are empty.

Now how do we turn a Young diagram into a permutation tableau? Easy, simply insert 1's and 0's into the boxes (those that aren't already occupied by 2's). But we have to follow two simple rules:

(I) Every column has to contain at least a single 1.

(II) No 0 can have a 1 anywhere to its north and a 1 anywhere to its west at the same time. [Throughout this paper we will use the cardinal directions instead of up, down, left, and right.]

Slightly restated, we can look at the second rule like this: if a 0 has a 1 to its north, then there must only be 0's to its west, and, alternatively, if it has a 1 to its west then there must only be 0's to its north. I also want to stress that when I say that a 1 lies "to the north of" a 0, then the 1 must share the same column, and, likewise, a 1 "to the west of" a 0 shares the same row. Figure 3 illustrates the type of situation that isn't allowed by rule (II).

Notice that rule (I) doesn't preclude the existence of a permutation tableau without any columns at all. In fact, for any n, there exists a unique permutation

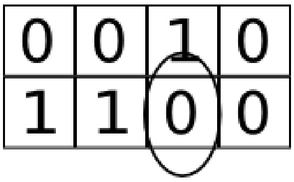

Figure 3: Forbidden by rule (II), as the circled 0 has a 1 to its north and a 1 to its west.

tableau with n rows and no columns (this will turn out to correspond to the identity permutation on n elements). On the other hand, rule (I) does imply that there are no diagrams with columns and no rows. Incidentally, rule (I) implies that there is a unique permutation tableau with $n-1$ columns and 1 row (it contains only 1's and we will see that it corresponds to the unique permutation on n elements that contains a single "nonwexbot").

What's a nonwexbot? A *weak excedance* for a permutation π is a pair $(i, \pi(i))$ of such that $i \leq \pi(i)$. Under this terminology, we call i a *weak excedance bottom* or *wexbot* and $\pi(i)$ a *weak excedance top* or *wextop*. A pair $(j, \pi(j))$ such that $j > \pi(j)$ is a *non-weak excedance* with j a *nonwexbot* and $\pi(j)$ a *nowextop*. While we're at it, let's set up one more concept that will help us create the bijection between permutation tableaux and permutations. A *primary path* is the unique path along the edges of a permutation tableau from the upper right-hand corner to the lower left hand-corner such that it keeps all the entries with 0's and 1's to its north and west. We could also think of it as the path that divides the 0's and 1's from the 2's. Note that if our permutation tableau is of size $m \times n$, then the primary path is always of length $m+n$. Figure 4 shows a permutation tableau (we'll call it $\mathbf{d_0}$ for future reference) with the primary path highlighted.

Now we will proceed to establish a bijection between these permutation tableaux and permutations by first defining a map ϕ from permutation tableaux to permutations and then a map θ from permutations and permutation tableaux and, finally, showing that the two maps are inverses of each other.

1.2 The Mapping ϕ

First, we'll describe the way that the permutation tableaux are mapped to permutation groups. It's tempting to jump to the conclusion that an $m \times n$ permutation tableau gets mapped to an element of S_m or S_n, but actually it gets mapped to S_{m+n}. From another angle, if a tableau's primary path is of length p then the

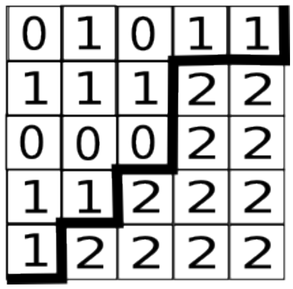

Figure 4: Permutation tableau $\mathbf{d_0}$. (The primary path is in bold.)

tableau is mapped to an element of S_p. Hence, tableaux of size 5×0, 4×1, 3×2, 2×3, and 1×4 are all mapped to elements of S_5.

So let's say that we're confronted with a tableau $\mathbf{d}$ that contains a primary path of length p. We find the tableau's image under ϕ by first labeling the edges of the path with the numbers from 1 to p starting with the northeastern edge and ending with the southwestern. An example using the tableau $\mathbf{d_0}$ appears in Figure 5; the squares with 2's have been removed so that the labeled path can be easily seen.

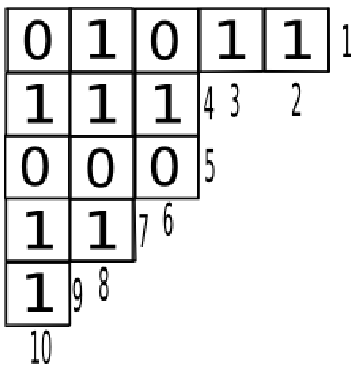

Figure 5: Labeled primary path.

Now we establish a simple method to determine where a number along the path gets mapped by the permutation $\phi(\mathbf{d})$. If a number k is associated with a vertical edge in the primary path, we consider the row in $\mathbf{d}$ that is directly west. Specifically, we look to see if any 1's are in the row. If not, then we say that k is fixed by $\phi(\mathbf{d})$ and it will be mapped to itself. On the other hand, if there are 1's in the row, then we imagine a path beginning at k and ending at the westmost 1 in the row. Then the path turns south and proceeds to alternate between east and south whenever it hits a 1. When the path hits an edge of the primary path once again, we say that the number associated with that edge is the image of k under $\phi(\mathbf{d})$. It follows that any number associated with a vertical edge of the primary path will be a wexbot in $\phi(\mathbf{d})$. Figure 6 gives an example of how to use the algorithm to find the images of a few of the wexbots in $\phi(\mathbf{d_0})$.

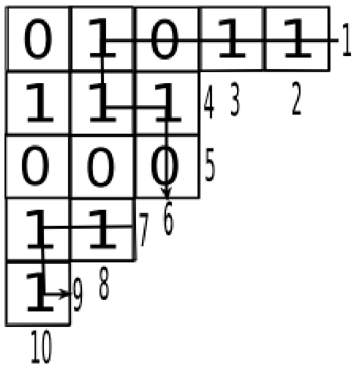

Figure 6: We see that $\phi(\mathbf{d_0})$ maps $1 \to 6$, $7 \to 9$, and $5 \to 5$.

Now suppose k is located on a horizontal edge of the primary path. We follow a similar procedure, although because of rule (I) we know a 1 will lie to the north of the edge associated with k. So imagine a path moving northward from k that only stops once it hits the northmost 1 in this column. Then the path turns east and continues turning east and south whenever it hits a 1. When the path reaches an edge of the primary path we say that the number associated with that edge is the image of k under $\phi(\mathbf{d})$. It follows that any number associated with horizontal edges of the primary path will be a nonwexbot in $\phi(\mathbf{d})$. In Figure 7, we can see the procedure carried out with a few of the nonwexbots in $\phi(\mathbf{d_0})$.

It will be useful for our purposes to prove that the south-east paths that are involved in defining $\phi(\mathbf{d})$ can neither "merge" nor "cross" each other. In order to be precise, it will be helpful to think of the path as an alternating sequence of vertices and edges. We will associate a vertex with each square in the permutation diagram and an edge with the shared edge of two adjoining squares. Each square, and therefore vertex, can be associated with a pair (j, k) in which j is a wexbot in $\phi(\mathbf{d})$ and k is a nonwexbot in $\phi(\mathbf{d})$.

Lemma 1.1 *The south-east paths defined above never share an edge and if two paths share a vertex then they don't intersect, they only touch (see Figure 8).*

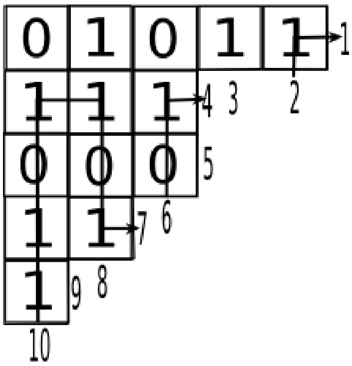

Figure 7: We can see that $\phi(\mathbf{d_0})$ maps $2 \to 1$, $6 \to 4$, and $10 \to 7$.

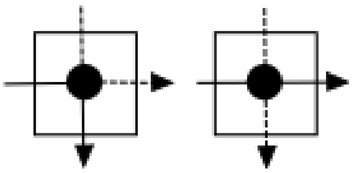

Figure 8: Two paths touch on the left and intersect on the right.

Proof Suppose a south-east path begins at some box (j,k). The box (j,k) must either contain a northmost 1 or a westmost 1. If (j,k) contains a northmost 1 then the path initially turns east and another south-east path can only merge with it at (j,k) if it enters (j,k) from the north which is impossible by the construction. Similarly, if (j,k) contains a westmost 1 then the path initially turns south and another south-east path can only merge with it at (j,k) if it enters (j,k) from the west which would also lead to a contradiction.

The only way that two south-east paths can share an edge is if they first meet at a vertex. We will show that any two paths that meet at a vertex will "touch" in the manner illustrated in Figure 8, and, therefore, they can't ever share an edge. Suppose south-east paths meet at a vertex (j,k). Since they aren't already merged, it must be the case that one path is approaching from the north and the other path is approaching from the west. This means that there must be a 1 to the west of (j,k) and a 1 to the north of (j,k). This implies, by rule (II), that the square (j,k) is occupied by a 1. Hence, the path approaching from the north will turn east and the path approaching from the west will turn south once they meet at (j,k). Therefore, the two paths will meet in the way we claimed above. ∎

Since $\phi(\mathbf{d})$ is a mapping from the finite set of elements to itself, we only need to make certain that the this mapping is injective in order to ascertain that it is a permutation.

Corollary 1.2 *For any permutation tableau* $\mathbf{d}$*, the mapping* $\phi(\mathbf{d})$ *is a permutation.*

Proof Lemma 1.1 tells us that $\phi(\mathbf{d})$ cannot map any two numbers to the same number. Hence $\phi(\mathbf{d})$ is injective. ∎

1.3 The Mapping θ

Now we will define a mapping θ from permutations to permutation tableaux. The initial goal will be to simply show that the image of a permutation under θ follows the rules laid out for permutation tableau. We will then proceed to prove that the image of $\phi(\mathbf{d})$ under θ is $\mathbf{d}$.

Given a permutation π, we arrive at the size of $\theta(\pi)$ by noting how many wexbots and how many nonwexbots are in π. The number of columns will correspond to the number of nonwexbots and the number of rows will correspond to the number of wexbots. Now we start with the rightmost column and assign 0's, 1's, and 2's to the boxes in each column by iterating the following procedure:

1. Note the number n_k of wexbots less than the nonwexbot k associated with the column (there must be at least one), and leave the northmost n_k squares empty while filling the remaining squares with 2's.

2. Note which edge is labeled with the image of the nonwexbot associated with the column. Insert a 1 in the box that will make $\phi(\theta(\pi))(k) = \pi(k)$. (This won't be a problem with the first column, simply place a 1 in box $(\pi(k), k)$, and we'll address the feasibility of carrying out this command for the remaining columns shortly.) Also place 0's in every box above the 1 just inserted.

3. Now we need to fill in any remaining empty squares below the 1. For the square (j, k), insert a 1 if $\phi(\theta(\pi))(j) \neq \pi(j)$ and a 0 if $\phi(\theta(\pi))(j) = \pi(j)$. If there are currently no 1's to the west of j then we will say, even if there aren't any 0's in the squares yet, that $\phi(\theta(\pi))(j) = \pi(j)$.

We continue with the exact same procedure for each subsequent column, but there is a question about whether we can always carry out the second step in our procedure. The following lemma addresses this issue.

Lemma 1.3 *For any nonwexbot* k *in a permutation* π*, a 1 can be inserted into a square* (j, k) *in* $\theta(\pi)$ *so that* $\phi(\theta(\pi))(k) = \pi(k)$ *(as long as 0's are inserted in all of the squares to the north of* (j, k)*).*

Proof We will prove this through induction, with the basis step already established by the parenthetical remark in step (2) of the procedure listed above. Now suppose the lemma holds for every column to the west of the column above the nonwexbot k. Locate the edge of the primary path labeled by $\pi(k)$. Trace a west-north (or north-west) path from $\pi(k)$ and insert a 1 at the point where the path intersects the column north of k. The only concern then is to show that the path doesn't hit the northmost edge of the tableau before reaching that column. But suppose it does. Then the path turned north after hitting the northernmost 1 in some column to the east of the desired column. Let k' be the nonwexbot associated with that column. Then $\phi(\theta(\pi))(k') = \pi(k)$ which contradicts our inductive hypothesis. ∎

This leads to the following corollary:

Corollary 1.4 *For any permutation π, $\phi(\theta(\pi))(k) = \pi(k)$ for all nonwexbots k in π.*

Our procedure, then, for the construction of $\theta(\pi)$ can be carried out and it only remains to establish that the resulting diagram is a permutation tableau, i.e. show that it follows rules (I) and (II).

Lemma 1.5 *For any permutation π, $\theta(\pi)$ is a permutation tableau.*

Proof Step (2) of the procedure of θ assures us that the resulting tableau will satisfy rule (I). A 0 will be inserted in $\theta(\pi)$ for one of two reasons. Suppose a 0 is inserted in the square (j, k). Either it was inserted during step (1) in which case we are done, or it was inserted during step (2). If inserted during step (2) then $\phi(\theta(\pi))(j) = \pi(j)$. This will continue to be the case in subsequent columns so there won't be any 1's inserted to the west of our 0 unless inserted because of rule (1). But if a 1 is inserted to the west of (j, k), say in square (j, k') then $\phi(\theta(\pi))(k') = \pi(j) \neq \pi(k')$ which contradicts step (3). ∎

1.4 Proof of the Bijection

We now proceed to prove that there exists a bijection between permutation tableaux and permutations by showing that ϕ and θ are inverses of each other.

Lemma 1.6 *ϕ followed by θ is the identity map.*

Proof We take an arbitrary permutation tableau $\mathbf{d}$ and show that the image of $\phi(\mathbf{d})$ under θ is $\mathbf{d}$. The image of $\phi(\mathbf{d})$ will have the same dimensions and principle path as $\mathbf{d}$, since these are derived from the number and identity of the wexbots and nonwexbots of $\phi(\mathbf{d})$ and those are derived from the dimensions and principle path of $\mathbf{d}$.

We just need to determine if the 1's and 0's in $\theta(\phi(\mathbf{d}))$ are placed the same as those for $\mathbf{d}$. We will prove this inductively by starting at the southmost row and moving northward. It has already been established that the 2's will be placed appropriately since the principle path will be in the same location in both diagrams. Will the 1's be put in the appropriate locations? Suppose the row is labeled with the wexbot j, then the westmost 1 in bottom row of $\mathbf{d}$ will be placed in the square $(j, \phi(\mathbf{d})(j))$ (by definition of ϕ) and this means that there will be a 1 placed in the same location in $\theta(\phi(\mathbf{d}))$ per step (3) of θ. Rules (I) and (II) dictate that the squares to the east of (j,$\phi(\mathbf{d})$(j)) will be filled with 1's in both $\mathbf{d}$ and $\theta(\phi(\mathbf{d}))$. The squares to the left will be filled with 0's in $\mathbf{d}$ by construction, and we can be assured that the same will be the case with $\theta(\phi(\mathbf{d}))$ because of step (3) of θ and Corollary 1.4 (i.e. no more 1's will be inserted at step (2) in the construction process.)

Suppose all the rows of $\mathbf{d}$ are identical with the rows of $\theta(\phi(\mathbf{d}))$ up until the row labeled with k'. Since the lower rows are all identical, we can be assured that the westmost 1 in the k'-row in $\theta(\phi(\mathbf{d}))$ will appear in the same square as the westmost 1 in $\mathbf{d}$. This follows from step (3) in the construction of $\theta(\phi(\mathbf{d}))$. Suppose a 0 occupies some square (k', j) to the east of that westmost 1 in $\mathbf{d}$. This means that there must be a northmost 1 in the j-column which lies to the south of the square (k', j), and corollary 1.4 tells us that $\theta(\phi(\mathbf{d}))$ has the same northmost 1 in its j-column. Hence, per rule (1), it has a 0 in square (k', j). Now suppose a 1 occupies a square (k', j) to the east of the westmost 1 in the k'-row of $\mathbf{d}$. It should be clear that the construction of $\theta(\phi(\mathbf{d}))$ dictates that there be a 1 in its square (k', j) as well. ∎

And now we prove the other direction:

Lemma 1.7 *θ followed by ϕ is the identity map.*

Proof We must show that given an arbitrary permutation π, both π and the image of $\theta(\pi)$ under ϕ are identical. Corollary 1.4 establishes the identity for nonwexbottoms. We will establish the identity for wexbottoms through induction on the rows of $\theta(\pi)$. Rule (3) of the process for creating $\theta(\pi)$ assures us that the wexbot associated with the lowest row of θ (i.e. the largest wexbot in π) will have the same image under $\phi(\theta(\pi))$ as under π. Now suppose k is a wexbot and for any wexbot k' associated with a row beneath k, that $\phi(\theta(\pi))(k') = \pi(k')$. Then following step (3) of θ, 1's will be placed in the row until $\phi(\theta(\pi))(k) = \pi(k)$. Is it possible for that never to occur? We clearly don't run into that problem if k is fixed under π. If it isn't, then form a north-west path starting at the edge of the principal path that is associated with $\pi(k)$. If the north-west path intersects the row associated with k, then a 1 can be inserted in the square in which the intersection occurs we'll have $\phi(\theta(\pi))(k) = \pi(k)$. If the path doesn't hit the row then it must hit the western edge of the tableaux, indicating that it hit a westmost 1 in a row below the k-row, and this means that the wexbot associated with that row will get mapped to $\pi(k)$ which contradicts our inductive hypothesis. ∎

The preceding two lemmas establish the following:

Theorem 1.8 *There is a bijection between permutation tableaux and permutations.* ∎

2 What Permutation Tableaux Tell Us

Now that the bijection has been established, the question arises about the purpose of such a bijection. First, there is information available to us from a permutation tableau that isn't quite as immediately apparent when we look at the permutation in any of its other traditional forms. For example, we immediately know how many wexbots and nonwexbots are in the permutation. Moreso, it is not, using this characterization, difficult to get a sense of the number of permutations on m elements that have a given number of wexbots. For example, while it may not be terribly difficult to reason out that there is only one permutation on m elements that contains m wexbots (i.e. the identity permutation.) It is perhaps a little more difficult, without the aid of the permutation tableau, to see that there is only one permutation that contains only a single wexbot (see Figure 9).

1	1	1	1	1

Figure 9: Recall that (I) tells us that every column must contain a 1.

An even more difficult calculation, without the use of these diagrams, is the number of permutations that have $m-1$ wexbots. The following pictures, though, allow us to relatively easily see such permutations number $(2^1 - 1) + (2^2 - 1) + (2^3 - 1) + ... + (2^{m-1} - 1) = 2^m - (m + 1)$.

There are only 4 possibilities here for the placement of 2's, and it can be seen that every possible combination of 1's and 0's can fit into the remaining squares except for the all-0's combination. Hence $2^4 - 1$ possibilities when there are no 2's, $2^3 - 1$ possibilities when there is one 2, etc.

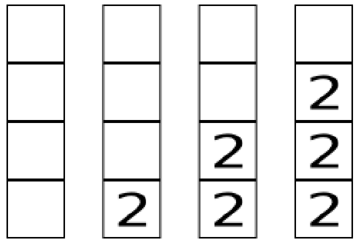

Figure 10: 4×1 permutation tableaux.

2.1 Permutation Statistics

We will describe some permutations statistics that give us a way to indicate the degree to which a given permutation is "mixed up". Before bringing up the permutation statistics let me introduce a different way of presenting a permutation which will allow us to understand the permutation statistics a little more easily. This presentation is referred to as a *chord diagram* and requires us to write the numbers involved around the edge of a circle and to use directed chords to indicate where the numbers are being mapped under the permutation. An example of a chord diagram can be seen in Figure 11.

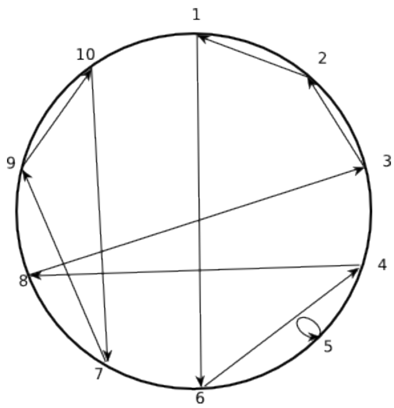

Figure 11: A chord diagram representation for $\phi(\mathbf{d})$.

Each element is associated with a chord, the outgoing edge from that vertex. From here on out, it should be understood that any reference to elements of the domain of a permutation "intersecting" or "not intersecting" means that their associated chords intersect or don't intersect.

Definition The domain of A_{EE} and C_{EE} is the set of wexbots of π where

$A_{EE}(i) =$ the set of nonintersecting wexbots greater than i,

$C_{EE}(i) =$ the set of intersecting wexbots less than i.

The domain of A_{NN}, A_{EN}, A_{NE} and C_{NN} is the set of nonwexbots of π and their outputs are defined as follows:

$A_{NN}(i) =$ the set of nonintersecting nonwexbots greater than i,

$A_{EN}(i) =$ the set of nonintersecting wexbots less than i,

$A_{NE}(i) =$ the set of nonintersecting wexbots greater than i,

$C_{NN}(i) =$ the set of intersecting nonwexbots greater than i.

Now we let $A_{EE}(\pi) =$ the sum of the cardinality of $A_{EE}(i)$ over every wexbot i in π, and use an analogous definition for the other six functions. Note that if j is in the image of $C_{EE}(i)$ then we will refer to the pair (i, j) as a crossing of type C_{EE}, and characterize crossings of type C_{NN} in the same manner.

This leads to some other useful aspects of the permutation tableaux, as illuminated by the following two theorems:

Theorem 2.1 *The number of 2's in the $n \times m$ permutation tableau T corresponds to $A_{NE}(\pi)$.*

Proof First we note that every 2 in T can be associated with a unique ordered pair consisting of the wexbot associated with the row it lies within and the nonwexbot associated with the column. Now it should be clear from the location of the 2 that the wexbot is larger than the nonwexbot and that they can't intersect because the wexbot is mapping to something equal to or larger than itself and the nonwexbot is mapping to something smaller than itself. The converse is fairly clear as well. Given any i and j that are related in that way, then the i is associated with some row and the j is associated with some column and by virtue of the fact that i is greater than j, then the intersection of the row and the column must contain a 2. ∎

Theorem 2.2 *There is a correspondence between 1's which aren't northmost in their respective columns and the elements of C_{EE} and C_{NN}.*

Proof The general idea is as follows, we will first show that the two paths in the permutation tableau associated with a crossing of type C_{EE} (alternatively, C_{NN}) intersect in a unique edge. Then we will show that these edges are in bijection with both the crossings and the non-northmost 1's in the diagram. This will establish what we want to prove.

We will prove the unique intersecting edge part of the proof for crossings of type C_{EE} noting that the proof for crossings of type C_{NN} proceeds in a completely analogous manner. Now let (i, j) be a crossing of type C_{EE} such that i maps to a_i and j maps to a_j. Now the path from i to a_i and the path from j to a_j (as dictated by ϕ) must intersect at some point since $i < j \leq a_i < a_j$ (see Figure 12 if you're unsure about this.)

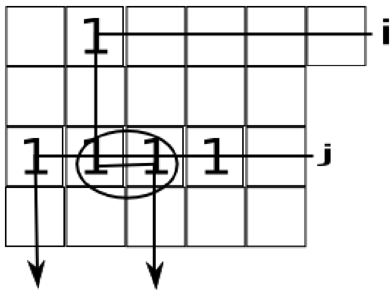

Figure 12: Crossing of the type C_{EE}.

Must the two paths intersect in an edge? The intersection must occur while the path from i to a_i is on its east-south section. The west-ward section of the path from j to a_j must intersect this east-south portion of the other path, otherwise a_j couldn't be larger than a_i. If this part of the path doesn't intersect in at least one edge then it must be traveling southward at the time and there must not be a 1 at the point of intersection. There must be a 1 to the north of that point of intersection due to the fact that the i-path is traveling southward and there must be a 1 to the west of that point since since the j-path has yet to reach the west-most 1 in its row. But this is a contradiction, so there must be a 1 at the point of intersection, and, therefore, the two paths must share at least one edge.

Now we need to establish that the paths share a unique edge. Suppose not. Then they must intersect while both paths are on their south-eastern path, but if that occurs in an edge, then the paths will have the same destination which is impossible, hence they intersect at a unique edge. Furthermore, the 1 associated with the initial point of intersection can't be a north-most 1 since it was approached from above by the i-path.

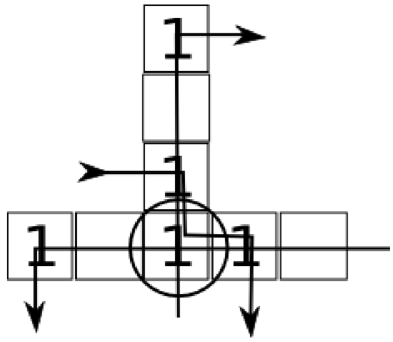

Figure 13: A contradiction: the east-south path can't both originate with a wexbot and a nonwexbot.

Alternatively, suppose that (i, j) is a crossing of type C_{NN}. We again can see that the two paths must intersect. In this case, the path from i to a_i must be on its east-south component when it crosses the northward component of the j-path. For the same reason as above, we see that the paths must intersect in an edge, and that they must only intersect in that one edge. Also, the 1 that marks the southern border of the edge of intersection clearly can't be the north-most 1 in its column.

So we choose to associate these unique edges with the crossings of type C_{EE} and type C_{NN}. Furthermore, we will look at these unique edges and if the edge is vertical then we will associate the southern 1 with the crossing and if the edge is horizontal then we will associate the western 1 with the crossing. So we must ask whether any 1's will be used twice in creating these associations. On examination, this couldn't be the case because it would mean that two of the paths would be shared and thus two numbers would be mapped to the same number under π which is a contradiction. This situation is made clear in Figure 13.

So we have an injective map between crossings and non-northmost 1's, now we just need to show that the mapping is also surjective. Let x be a non-northmost 1 and suppose that it isn't associated with a crossing of type C_{EE}. [We must show that it is associated with a crossing of type C_{NN}.] We start by tracing a path northward from x and turning west at the first 1 that we encounter. We then turn north at the next 1 and keep following that north-west pattern until we hit a northmost 1 when moving westward, which must happen sooner or later. Why? The only way that this couldn't happen is if we hit a 1 from the south and there isn't a 1 to the west, but if that were to happen, we can see that we will have hit

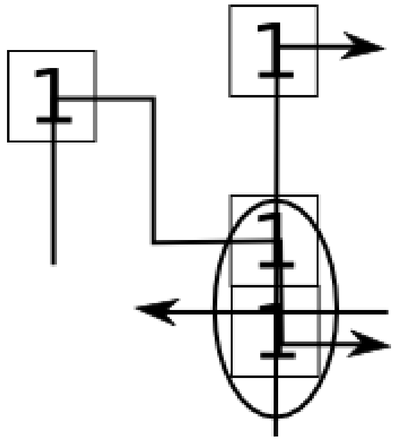

Figure 14: We can see the formation of a type C_{NN} crossing as well as how a type C_{EE} crossing would be formed if the east-south path originated from a west-most 1.

the westmost 1 in a row, which means that the number associated with that row has a path that runs through the 1 at which we started out. But this contradicts our hypothesis, because it means that the 1 is associated with a crossing of type C_{EE}. We then travel south from that 1 and see that the path that we were on was the path from the number associated with the column that we're now within to its image. It can be seen (look at Figure 14) that the number associated with this column and the number associated with the column in which we started out form a crossing of type C_{NN} which is what we needed to show. Thus our bijection is established. ∎

3 Permutation Tableaux and Descents

A second bijection between permutation tableaux and permutations will be established which will allow us to make a connection between descents (defined below) and wexbots, and from this a connection between descents and the 1's, 0's, and 2's of the permutation tableau. For ease of exposition, we will be representing a given permutation π by its image. That is to say, 43157862 is the permutation that maps 1 to 4, 2 to 3, 3 to 1, etc. We will be abusing the terminology and referring to the sequence of numbers that represents the image of π as π itself, and referring to numbers in the sequence as elements of π.

In order to define a descent we first note that a *permutation pattern* is a series of numbers which represent the relative order of numbers after a permutation has been applied to them. A dash between numbers in a permutation pattern indicates that the numbers don't have to be adjacent and the lack of a dash means that they do have to be adjacent. These permutation patterns will be treated as functions that map a permutation to the number that represents the number of times that the pattern occurs in the permutation. Some examples are in order. Let's say that $\pi = 43157862$, then $(21\text{-}3)\pi = 8$ since π contains the subsequences 435, 437, 438, 436, 315, 317, 318, and 316. All of which possess the same relative order (along with the adjacency of) as 21-3.

A *descent* is a pair of two adjacent numbers in a permutation such that the second number is less than the first. The permutation pattern 456213 contains two descents, 62 and 21. We will refer to the lesser number in a descent as a *descent bottom* and the greater number in a descent as a *descent top.* As with 2 in the preceding example, a number can both be a descent bottom and a descent top.

The particular permutation patterns on which we will be focusing are the 6 three-digit patterns that contain a single descent. These are (21-3), (31-2), (32-1), (1-32), (2-31), and (3-21). We will establish a correspondence between these permutation patterns and the permutation statistics described earlier. In order to do that, we need to establish a bijection and, before doing that, we'll define two more concepts that will arise when we define the mapping ψ. The *right-embracing number* of x in π, or REMBR(x), is the number of occurrences of the permutation pattern 2-31 in which x occupies the 2 place.

3.1 Description of the Permutation Operator ψ

Given a permutation $\pi \in S_n$, first form the following four sets:

$A =$ The descent bottoms in π.

$B =$ The descent tops in π.

$C =$ The non-descent bottoms in π (note that this can overlap with B.)

$D =$ The non-descent tops in π (note that this can overlap with A.)

Note the right-embracing number for each element of π. Now we form two new sets:

$A^* = \{i + 1 | i \in A\} \cup \{1\}$

$B^* = \{i - 1 | i \in B\} \cup \{n\}$

We will map the elements of A^* to the elements of B^* so that REMBR(x) for x in π equals the size of $C_{EE}(x)$ in the new permutation $\psi(\pi)$. It will be proven that this can always be done shortly, but, for the moment, I will simply take the feasibility for granted and show how one goes about doing it. Start with the largest element x of A^* and map it to the element y of B^* such that there are exactly REMBR(x) elements of B^* that are greater than or equal to x and less than y. Pick the next smallest number in A^* and continue the procedure until all of the elements of A^* are assigned an element from B^*. We will carry out a similar procedure with the sets C and D. First form two new sets:

$C^* = \{i + 1 | i \in C - \{n\}\}$

$D^* = \{i - 1 | i \in D - \{1\}\}$

This time the elements of C^* are mapped to the elements in D^* so that REMBR(x) from x in π equals the size of $C_{NN}(x)$ in the new permutation $\psi(\pi)$. Take the smallest element x of C^* and map it to the element y of D^* such that the number of elements greater than y and less than or equal to x equals REMBR(x). Repeat this procedure with the next largest element (just don't map to an element if it's the image of somebody else) of C^* until every element of C^* has been assigned to an element of D^*.

At this point the union of the relations created thus far will be a permutation $\psi(\pi)$ on the same number of elements as π. A lot of questions are raised by this procedure. How do we know that it can be carried out? How do we know that the resulting set of relations is a permutation? I will first carry out a "proof" by example and then present the actual proof that this mapping works and is, in fact, a bijection. Let's try to the procedure with the permutation $\pi = 731986452$. The sets A, B, C, and D are as follows:

$A = \{3, 1, 8, 6, 4, 2\}$

$B = \{7, 3, 9, 8.6, 5\}$

$C = \{7, 9, 4\}$

$D = \{1, 4, 2\}$

Note that $REMBR(7) = REMBR(3) = REMBR(4) = 1$ and 0 for the remaining numbers. We form the sets A^* and B^*:

$A^* = \{1, 2, 3, 4, 5, 7, 9\}$

$B^* = \{2, 4, 5, 6, 7, 8, 9\}$

Following the procedure above we get $\psi(\pi)(9) = 9$, $\psi(\pi)(7) = 8$, $\psi(\pi)(5) = 5$, $\psi(\pi)(4) = 6$, $\psi(\pi)(3) = 7$, $\psi(\pi)(2) = 2$, $\psi(\pi)(1) = 4$. The ease with which this procedure works practically begs for the proof that will be shortly provided. The remaining construction is quickly accomplished:

$C^* = \{6, 8\}$

$D^* = \{1, 3\}$

So we find that $\psi(\pi)(6) = 1$ and $\psi(\pi)(8) = 3$, and our final permutation is $\psi(731986452) = 427651839$. We see that the image of our original permutation is indeed a permutation. Now let's see why that occurs and prove that the mapping actually gives us a bijection. Before we present our theorem, however, we will prove a lemma that will prove crucial. We will define DESTOP(π) as the set of descent tops for π and DESBOT(π) as the set of descent bottoms for π. The following lemma was originally established through the proof of Theorem 4 in [4], but I carried out an independent proof for this paper.

Lemma 3.1 *A permutation π is completely determined by DESTOP(π), DESBOT(π), and $\{REMBR(x)|x \in \pi\}$.*

Proof First observe that any permutation consists of a series of decreasing and increasing intervals. The set $A = DESTOP(\pi) - DESBOT(\pi)$ consists of the elements that begin decreasing intervals and the set $B = DESBOT(\pi) - DESTOP(\pi)$ consists of the elements that end decreasing intervals. It follows that sets A and B must be the same size. Suppose $[a, b]$ is a decreasing interval, then REMBR(x) for $x \in \pi$ increases by 1 if and only if $a > x > b$ and $[a, b]$ lies to the right of x in π. Now suppose REMBR(x) $= m$, then there must be m intervals of the type described above that lie to the right of x. Furthermore, suppose $x \in DESTOP(\pi) \cup DESBOT(\pi)$, then it follows that x must lie in the $m + 1$st descending interval (counting from right to left) that could potentially contain x. On the other hand, suppose x is neither an element of DESTOP(π) nor an element of DESBOT(π), then x is located on an increasing interval, and, furthermore, must lie on the first increasing interval that can potentially contain x that lies to the left of the first m decreasing intervals (again counting right to left) such that $a > x > b$. The point being, once we determine the bounds of the decreasing intervals in our permutation, we can then use the right-embracing numbers to determine where every other element of the permutation lies.

Now we demonstrate how to determine the bounds of the decreasing intervals. Let $B* = \{x \in B \mid x < y \; \forall \; y \in A\}$ and label the elements so that $b_m < \cdots < b_2 < b_1$. Now we note that there must be exactly REMBR(a_1) elements of B* that lie to the right of b_1. We can repeat the process with a_2, noting that REMBR(a_2)

elements of B^* (not including a_1) must lie to the right of a_2. It should be clear that the ordering of the elements of B^* is completely determined in this manner (an example will follow). Label the elements of $A \cup (B - B^*)$ so that $c_j < \cdots < c_2 < c_1$. Due to the definition of B^*, c_1 must be an element of DESTOP(π). We note that c_1 must be placed to the left of the rightmost REMBR(c_1)+1 elements of B^* (the +1 because c_1 must be associated with a member of DESBOT(π) and that can't be part of one of the descents that will embrace c_1 on the right). If c_2 is also a member of DESTOP(π) then we can place it in the ordering just as we did with c_1, the only difference being that we can't count the descent defined by c_1 and its bottom because it certainly won't embrace c_2. If c_2 is an element of DESBOT(π) then we can place it in the same way except that it will be placed to the left of the rightmost REMBR(c_2) elements of DESBOT(π) that aren't already linked with an element of DESTOP(π). Repeating this procedure with the remaining elements will allow us to determine the location of all the decreasing intervals for π and so, as noted above, we can then completely determine the permutation π. ∎

Here is a demonstration of the procedure outlined above.

Example Let $\pi = 452178396$. Then DESTOP(π)= $\{5, 2, 8, 9\}$ and DESBOT(π)= $\{2, 1, 3, 6\}$. We also note that REMBR(4) = REMBR(7) = 2, REMBR(5) = REMBR(8) = 1, and the right-embracing number is 0 for the remaining elements of π. Now we need to use this info to recreate π. First we note that $A = \{5, 8, 9\}$ and $B = \{1, 3, 6\}$. Then $B^* = \{1, 3\}$ and since both 1 and 3 have a right-embracing number of 0, it follows that their order must be "13". Now, since REMBR(5)=1 we get "513". Continuing to follow the procedure outlined above, we get "5136", "51836", and finally "518396". So the descending intervals of π are [5,1], [8,3], and [9,6]. Now we can place the remaining elements of π according to their right-embracing numbers. This gives us "5218396", "45218396", and finally we end up with "452178396" which is what we wanted to end up with.

Now we are able to use the lemma to establish the second bijection.

Theorem 3.2 *The mapping ψ is an operator on the permutations.*

Proof I will start by demonstrating that we will always be able to assign images to a wexbot x in the domain of $\psi(\pi)$ so that the size of of $C_{EE}(x)$ is always the same size as REMBR(x) in π. First look at the largest wexbot, let's call it x_1, in $\psi(\pi)$ and observe that there are at least REMBR(x_1) descent tops in π that are greater than x_1. Due to the way that B^* was defined, this implies that there are at least REMBR(x_1) + 1 numbers in B^* that are greater than or equal to x_1. This in turn tells us that we can certainly map x_1 to an element of B^* so that $C_{EE}(x_1)$ in $\psi(\pi)$ equals REMBR(x_1) in π.

Now we show inductively that a similar situation will hold for all the wexbots in $\psi(\pi)$. Suppose that for the first k wexbots (moving from greatest to smallest) in $\psi(\pi)$ we can map them to an element in the domain so that the number of crossings of C_{EE} are the same as the right-embracing number in π. [We must show that x_{k+1} can be mapped in such a way.] As above, we notice that there are at least REMBR$(x_{k+1}) + 1$ elements of B^* that are greater than or equal to x_{k+1}. Unfortunately, k of those elements must already be the image of elements, x_1 through x_k, that are larger than x_{k+1} and so they can't contribute to the size of $C_{EE}(x_{k+1})$. On the other hand, recall that the wexbots for $\psi(\pi)$ are produced by adding 1 to the descent bottoms of π and so the first k wexbots all correspond to descent bottoms that are larger than x_{k+1}. Now there are, of course, descent tops that are larger still that go with these descent bottoms and none of these descents can possibly embrace x_{k+1}. This shows us that REMBR(x_{k+1}) is no greater than the number of descent tops to the right of x_{k+1} in π minus k. Therefore we will be able to find an element of B^* such that there are REMBR(x_{k+1}) numbers smaller than it and greater than or equal to x_{k+1}.

The procedure for assigning elements of C^* to elements of D^* works for analogous reasons and can be shown inductively in the same manner. We also know that the image of an element of S_n under ψ is also an element of S_n so now we simply need to prove injectivity in order to establish the permutation.

Suppose $\psi(\pi_1) = \psi(\pi_2)$ for $\pi_1, \pi_2 \in S_n$. [We must show that $\pi_1 = \pi_2$.] It must be the case that both permutations share descent tops, descent bottoms, and right-embracing numbers and so by Lemma 3.1 it follows that $\pi_1 = \pi_2$. ∎

Note that the bijection implicitly highlights a symmetry between descents and weak-excedances. In fact, the collection of permutations with k weak excedances are in bijection with the collection of permutations with $k-1$ descents. This gives us a nice menu of approaches for dealing with counting problems linked to the number of weak excedances or descents.

4 Concluding Remarks

Any permutation can be represented using a permutation tableau or mapped to a permutation with the number of descents linked to the number of wexbots in the original permutation. In trying to count permutation tableaux, one is naturally drawn to the problem of counting the number of permutations on n elements with a given number of wexbots, and we saw, above, how permutation tableaux can simplify this problem in some situations. One is also quickly confronted with the apparent symmetry between the number of permutations with $n-k+1$ weak excedances and k weak excedances, e.g. there is exactly one permutation on n

elements with n weak excedances and only one, as well, with a single weak excedance. Both of the bijections established above give us different, possibly more fruitful ways, of approaching this and other similar problems. I've not established with certainty whether either or both of these problems are open, but I've not encountered any solutions in the literature that I've surveyed.

References

[1] E. Steingrimsson, L.K. Williams, Permutation tableaux and permutation patterns, Journal of Combinatorial Theory, Series A 114 (2007) 211-234.

[2] A. Postnikov, Webs in totally positive Grassmann cells, in preparation.

[3] L. Williams, Shelling totally nonnegative flag varieties, preprint, math.RT/0509129.

[4] R.J. Clarke, E. Steingrimsson, J. Zeng, New Euler-Mahonian statistics on permutations and words, Adv. in Appl. Math. 18 (1997) 237-270.