



ELSEVIER

Available at
www.ElsevierMathematics.com
POWERED BY SCIENCE @ DIRECT®

Journal of Combinatorial Theory, Series A 105 (2004) 207–219

Journal of
Combinatorial
Theory
Series A

<http://www.elsevier.com/locate/jcta>

Bijections for refined restricted permutations

Sergi Elizalde and Igor Pak

Department of Mathematics, MIT, Cambridge, MA, 02139, USA

Received 11 January 2003

Abstract

We present a bijection between 321- and 132-avoiding permutations that preserves the number of fixed points and the number of excedances. This gives a simple combinatorial proof of recent results of Robertson et al. (Ann. Combin. 6 (2003) 427), and Elizalde (Proc. FPSAC 2003). We also show that our bijection preserves additional statistics, which extends the previous results.

© 2003 Elsevier Inc. All rights reserved.

Keywords: Restricted permutations; Bijection; Pattern avoidance; Permutation statistics

1. Introduction

The subject of *pattern avoiding permutations*, also called *restricted permutations*, has blossomed in the past decade. A number of enumerative results have been proved, new bijections found, and connections to other fields established. Despite recent progress, the so-called Stanley–Wilf conjecture giving an exponential upper bound on the number of pattern avoiding permutations remains open, and much of the ongoing research is related to the conjecture.

An unexpected recent result of Robertson et al. [10] gives a new and exciting extension to what is now regarded as a classical result that the number of 321-avoiding permutations equals the number of 132-avoiding permutations. They show that one can “refine” this result by taking into account the number of fixed points in a permutation. In fact, they study all six patterns in \mathcal{S}_3 which produce different “refined” statistics, with the above-mentioned result having a highly nontrivial and technically involved proof. The story continued in a recent paper of Elizalde [4]

E-mail addresses: sergi@math.mit.edu (S. Elizalde), pak@math.mit.edu (I. Pak).

where an additional statistic, “the number of excedances”, was added. The proof uses some nontrivial generating function machinery and is also quite involved.

In this paper we present a bijective proof of the “refined” results on 321- and 132-avoiding permutations, resolving the problem which was left open in [10,4]. In fact, our bijection is a composition of two (slightly modified) known bijections into Dyck paths, and the result follows from a new analysis of these bijections. The Robinson–Schensted–Knuth (RSK) correspondence is a part of one of them, and the difficulty of the analysis stems from the complexity of this celebrated correspondence. As a new application of our bijections, we show that the length of the longest increasing subsequence in 321-avoiding permutations corresponds to a certain statistic (that we call *rank*) in 132-avoiding permutations, which further refines the previous results. We also apply our bijections to “refined restricted involutions” (see Section 6).

Let n, m be two positive integers with $m \leq n$, and let $\sigma = (\sigma(1), \sigma(2), \dots, \sigma(n)) \in \mathcal{S}_n$ and $\pi = (\pi(1), \pi(2), \dots, \pi(m)) \in \mathcal{S}_m$. We say that σ *contains* π if there exist indices $i_1 < i_2 < \dots < i_m$ such that $(\sigma(i_1), \sigma(i_2), \dots, \sigma(i_m))$ is in the same relative order as $(\pi(1), \pi(2), \dots, \pi(m))$. If σ does not contain π , we say that σ is π -*avoiding*. For example, if $\pi = 132$, then $\sigma = (2, 4, 5, 3, 1)$ contains 132, because the subsequence $(\sigma(1), \sigma(3), \sigma(4)) = (2, 5, 3)$ has the same relative order as $(1, 3, 2)$. However, $\sigma = (4, 2, 3, 5, 1)$ is 132-avoiding.

We say that i is a *fixed point* of a permutation σ if $\sigma(i) = i$. Similarly, i is an *excedance* of σ if $\sigma(i) > i$. Denote by $\text{fp}(\sigma)$ and $\text{exc}(\sigma)$ the number of fixed points and the number of excedances of σ , respectively.

Denote by $\mathcal{S}_n(\pi)$ the set of π -avoiding permutations in \mathcal{S}_n . For the case of patterns of length 3, it is known [6] that regardless of the pattern $\pi \in \mathcal{S}_3$, $|\mathcal{S}_n(\pi)| = C_n = \frac{1}{n+1} \binom{2n}{n}$, the n th Catalan number. While the equalities $|\mathcal{S}_n(132)| = |\mathcal{S}_n(231)| = |\mathcal{S}_n(312)| = |\mathcal{S}_n(213)|$ and $|\mathcal{S}_n(321)| = |\mathcal{S}_n(123)|$ are straightforward, the equality $|\mathcal{S}_n(321)| = |\mathcal{S}_n(132)|$ is more difficult to establish. Bijective proofs of this fact are given in [7,9,12,14]. However, none of these bijections preserves either of the statistics $\text{fp}(\cdot)$ or $\text{exc}(\cdot)$.

Theorem 1 (Robertson et al. [10], Elizalde [4]). *The number of 321-avoiding permutations $\sigma \in \mathcal{S}_n$ with $\text{fp}(\sigma) = i$ and $\text{exc}(\sigma) = j$ equals the number of 132-avoiding permutations $\sigma \in \mathcal{S}_n$ with $\text{fp}(\sigma) = i$ and $\text{exc}(\sigma) = j$, for any $0 \leq i, j \leq n$.*

A special case of the theorem, which ignores the number of excedances, was given in [10]. In full, the theorem was shown in [4]. As we mentioned above, both proofs are non-bijective and technically involved. The main result of this paper is a bijective proof of the following extension of Theorem 1.

Let $\text{lis}(\sigma)$ be the length of the *longest increasing subsequence* of σ , i.e., the largest m for which there exist indices $i_1 < i_2 < \dots < i_m$ such that $\sigma(i_1) < \sigma(i_2) < \dots < \sigma(i_m)$. Define the *rank* of σ , denoted $\text{rank}(\sigma)$, to be the largest k such that $\sigma(i) > k$ for all $i \leq k$. For example, if $\sigma = 63528174$, then $\text{fp}(\sigma) = 1$, $\text{exc}(\sigma) = 4$, $\text{lis}(\sigma) = 3$ and $\text{rank}(\sigma) = 2$.

Theorem 2. *The number of 321-avoiding permutations $\sigma \in \mathcal{S}_n$ with $\text{fp}(\sigma) = i$, $\text{exc}(\sigma) = j$ and $\text{lis}(\sigma) = k$ equals the number of 132-avoiding permutations $\sigma \in \mathcal{S}_n$ with $\text{fp}(\sigma) = i$, $\text{exc}(\sigma) = j$ and $\text{rank}(\sigma) = n - k$, for any $0 \leq i, j, k \leq n$.*

To prove this theorem, we establish a bijection Θ between $\mathcal{S}_n(321)$ and $\mathcal{S}_n(132)$, which respects the statistics as above. While Θ is not hard to define, its analysis is less straightforward and will occupy much of the paper.

The rest of the paper is structured as follows. In Section 2 we define Dyck paths and several new statistics on them. The description of the main bijection is done in Section 3, and is divided into two parts. First we give a bijection from 321-avoiding permutations to Dyck paths, and then another one from Dyck paths to 132-avoiding permutations. In Section 4 we establish properties of these bijections which imply Theorem 2. Section 5 contains proofs of two technical lemmas. We conclude with extensions of our results to refined restricted involutions, and other applications.

Let us mention here that whenever possible we refer to the celebrated monograph [13] rather than to the original source. The interested reader is advised to consult [13] for the details, history, and further references on the subject.

2. Statistics on Dyck paths

Recall that a *Dyck path* of length $2n$ is a lattice path in \mathbb{Z}^2 between $(0, 0)$ and $(2n, 0)$ consisting of up-steps $(1, 1)$ and down-steps $(1, -1)$ which never goes below the x -axis. Sometimes it will be convenient to encode each up-step by a letter u and each down-step by d , obtaining an encoding of the Dyck path as a *Dyck word*. We shall denote by \mathcal{D}_n the set of Dyck paths of length $2n$, and by $\mathcal{D} = \bigcup_{n \geq 0} \mathcal{D}_n$ the class of all Dyck paths.

For any $D \in \mathcal{D}$, we define a *tunnel* of D to be a horizontal segment between two lattice points of D that intersects D only in these two points, and stays always below D . Tunnels are in obvious one-to-one correspondence with decompositions of the Dyck word $D = AuBdC$, where $B \in \mathcal{D}$ (no restrictions on A and C). In the decomposition, the tunnel is the segment that goes from the beginning of u to the end of d . If $D \in \mathcal{D}_n$, then D has exactly n tunnels, since such a decomposition can be given for each up-step of D .

A tunnel of $D \in \mathcal{D}_n$ is called a *centered tunnel* if the x coordinate of its midpoint (as a segment) is n , that is, the tunnel is centered with respect to the vertical line through the middle of D . In terms of the decomposition of the Dyck word $D = AuBdC$, this is equivalent to A and C having the same length $|A| = |C|$. Alternatively, this can be taken as a definition of centered tunnel. Throughout the paper we denote by $\text{ct}(D)$ the number of centered tunnels of D .

A tunnel of $D \in \mathcal{D}_n$ is called a *right tunnel* if the x coordinate of its midpoint is strictly greater than n , that is, the midpoint of the tunnel is to the right of the vertical line through the middle of D . In terms of the decomposition $D = AuBdC$, this is equivalent to saying that $|A| > |C|$. Denote by $\text{rt}(D)$ the number of right tunnels of D .




Fig. 1. One centered and four right tunnels.

In Fig. 1, there is one centered tunnel drawn with a solid line, and four right tunnels drawn with dotted lines. Similarly, a tunnel is called a *left tunnel* if the x coordinate of its midpoint is strictly less than n . Denote by $lt(D)$ the number of left tunnels of D . Clearly, $lt(D) + rt(D) + ct(D) = n$ for any $D \in \mathcal{D}_n$.

We will distinguish between right tunnels of $D \in \mathcal{D}_n$ that are entirely contained in the half plane $x \geq n$ and those that cross the vertical line $x = n$. These will be called *right-side tunnels* and *right-across tunnels*, respectively. In terms of Dyck words, a decomposition $D = AuBdC$ corresponds to a right-side tunnel if $|A| \geq n$, and to a right-across tunnel if $|C| < |A| < n$. In Fig. 1 there are three right-side tunnels and one right-across tunnel. *Left-side tunnels* and *left-across tunnels* are defined analogously.

Finally, for any $D \in \mathcal{D}_n$, define $v(D)$ to be the height of the middle point of D , that is, the y coordinate of the intersection of the vertical line $x = n$ with the path. For the path in Fig. 1, $v(D) = 2$.

We say that i is an *antiexcedance* of σ if $\sigma(i) < i$. Sometimes it will be convenient to represent a permutation $\sigma \in \mathcal{S}_n$ as an $n \times n$ array with a cross on the squares $(i, \sigma(i))$. Note that fixed points, excedances, and antiexcedances correspond respectively to crosses on, strictly to the right, and strictly to the left of the main diagonal of the array.

3. Two bijections into Dyck paths

The bijection $\Theta : \mathcal{S}_n(321) \rightarrow \mathcal{S}_n(132)$ that we present will be the composition of two bijections, one from $\mathcal{S}_n(321)$ to \mathcal{D}_n , and another one from \mathcal{D}_n to $\mathcal{S}_n(132)$.

The first bijection $\Psi : \mathcal{S}_n(321) \rightarrow \mathcal{D}_n$ is defined in two steps. Given $\sigma \in \mathcal{S}_n(321)$, we start by applying the RSK correspondence to σ [13, Section 7.11] (see also [6]). This correspondence gives a bijection between the symmetric group \mathcal{S}_n and pairs (P, Q) of *standard Young tableaux* of the same shape $\lambda \vdash n$. For $\sigma \in \mathcal{S}_n(321)$ the algorithm is particularly easy because in this case the tableaux P and Q have at most two rows. The *insertion tableau* P is obtained by reading σ from left to right and, at each step, inserting $\sigma(i)$ to the partial tableau obtained so far. Assume that $\sigma(1), \dots, \sigma(i-1)$ have already been inserted. If $\sigma(i)$ is larger than all the elements on the first row of the current tableau, place $\sigma(i)$ at the end of the first row. Otherwise, let m be the leftmost element on the first row that is larger than $\sigma(i)$. Place $\sigma(i)$ in the square that m occupied, and place m at the end of the second row (in this case we say that $\sigma(i)$ *bumps* m). The *recording tableau* Q has the same shape as P and is obtained by placing i in the position of the square that was created at step i (when $\sigma(i)$ was inserted) in the construction of P , for all i from 1 to n . We write $RSK(\sigma) = (P, Q)$.




Fig. 2. Construction of the RSK correspondence $\text{RSK}(\sigma) = (P, Q)$ for $\sigma = (2, 3, 5, 1, 4, 6, 8, 7)$.

Now, the first half of the Dyck path $\Psi(\sigma)$ is obtained by adjoining, for i from 1 to n , an up-step if i is on the first row of P , and a down-step if it is on the second row. Let A be the corresponding word of u 's and d 's. Similarly, let B be the word obtained from Q in the same way. We define $\Psi(\sigma)$ to be the Dyck path obtained by the concatenation of the word A and the word B written backwards. For example, from the tableaux P and Q as in Fig. 2 we get the Dyck path shown in Fig. 1. The following proposition summarizes properties of this bijection Ψ :

Proposition 3. *The bijection $\Psi : \mathcal{S}_n(321) \rightarrow \mathcal{D}_n$ satisfies $\text{fp}(\sigma) = \text{ct}(\Psi(\sigma))$, $\text{exc}(\sigma) = \text{rt}(\Psi(\sigma))$, and $\text{lis}(\sigma) = \frac{1}{2}(n + v(\Psi(\sigma)))$, for all $\sigma \in \mathcal{S}_n(321)$.*

Suppose $\text{RSK}(\sigma) = (P, Q)$ for any $\sigma \in \mathcal{S}_n$. A fundamental and highly nontrivial property of the RSK correspondence is the *duality*: $\text{RSK}(\sigma^{-1}) = (Q, P)$ [13, Section 7.13]. The classical *Schensted's Theorem* states that $\text{lis}(\sigma)$ is equal to the length of the first row of the tableau P (and Q). Both results are used in the proof of Proposition 3.

Let us now define the second bijection $\Phi : \mathcal{S}_n(132) \rightarrow \mathcal{D}_n$ as follows. Any permutation $\sigma \in \mathcal{S}_n$ can be represented as an $n \times n$ array with crosses in positions $(i, \sigma(i))$. From this array of crosses, we obtain the *diagram* of σ as follows. For each cross, shade the cell containing it and the squares that are due south and due east of it. The diagram is the region that is left unshaded. It is shown in [8] that this gives a bijection between $\mathcal{S}_n(132)$ and Young diagrams that fit in the shape $(n - 1, n - 2, \dots, 1)$. Consider now the path determined by the border of the diagram of σ , that is, the path with *up* and *right* steps that goes from the lower-left corner to the upper-right corner of the array, leaving all the crosses to the right, and staying always as close to the diagonal connecting these two corners as possible. Define $\Phi(\sigma)$ to be the Dyck path obtained from this path by reading an up-step every time it goes up and a down-step every time it goes right. Since the path in the array does not go below the diagonal, $\Phi(\sigma)$ does not go below the x -axis.

The bijection Φ is essentially the same bijection between $\mathcal{S}_n(132)$ and \mathcal{D}_n given by Krattenthaler [7] (see also [5]), up to reflection of the path from a vertical line (Fig. 3).

Proposition 4. *The bijection $\Phi : \mathcal{S}_n(132) \rightarrow \mathcal{D}_n$ satisfies $\text{fp}(\sigma) = \text{ct}(\Phi(\sigma))$, $\text{exc}(\sigma) = \text{rt}(\Phi(\sigma))$, and $\text{rank}(\sigma) = \frac{1}{2}(n - v(\Phi(\sigma)))$, for all $\sigma \in \mathcal{S}_n(132)$.*




Fig. 3. The bijection $\Phi : (6, 7, 4, 3, 5, 2, 8, 1) \mapsto \text{uduuduududdduddd}$.




Fig. 4. A cross and the corresponding tunnel.

Proof. Let us show using the diagram representation that Φ maps fixed points to centered tunnels and excedances to right tunnels. To do that we define the inverse map $\Phi^{-1} : \mathcal{D}_n \rightarrow \mathcal{S}_n(132)$. Given a Dyck path $D \in \mathcal{D}_n$, the first step needed to reverse the above procedure is to transform D into a path U from the lower-left corner to the upper-right corner of an $n \times n$ array, not going below the diagonal connecting these two corners. Then, the squares to the left of this path form a diagram, and we can shade all the remaining squares. From this diagram, the permutation $\sigma \in \mathcal{S}_n(132)$ can be recovered as follows: row by row, put a cross in the leftmost shaded square such that there is exactly one cross in each column. Start from the top and continue downward until all crosses are placed.

For the proof of this proposition, instead of using $D = \Phi(\sigma)$, it will be convenient to consider the path U from the lower-left corner to the upper-right corner of the array of σ . We will talk about tunnels of U to refer to the corresponding tunnels of D under this trivial transformation.

Consider the arrangement of crosses of σ as defined earlier. We now show how to associate a unique tunnel of D to each cross of this array. Observe that given a cross in position (i, j) , U has a vertical step in row i and a horizontal step in column j . In D , these two steps correspond to steps u and d , respectively, so they determine a decomposition $D = AuBdC$ (see Fig. 4), and therefore a tunnel of D . According to whether the cross was to the left of, to the right of, or on the main diagonal, the associated tunnel will be, respectively, a left, right, or centered tunnel of D . Thus, fixed points give centered tunnels and excedances give right tunnels.

To prove the last equality of the proposition, notice that $\text{rank}(\sigma)$ is the largest m such that an $m \times m$ square fits in the upper-left corner of the diagram of σ . Therefore, the height of $\Phi(\sigma)$ at the middle is exactly $v(\Phi(\sigma)) = n - 2 \text{rank}(\sigma)$. \square

The main result of the paper follows now easily from these two propositions.

Proof of Theorem 2. Propositions 3 and 4 imply that $\Theta = \Phi^{-1} \circ \Psi$ is a bijection from $\mathcal{S}_n(321)$ to $\mathcal{S}_n(132)$ which satisfies $\text{fp}(\Theta(\sigma)) = \text{ct}(\Psi(\sigma)) = \text{fp}(\sigma)$, $\text{exc}(\Theta(\sigma)) = \text{rt}(\Psi(\sigma)) = \text{exc}(\sigma)$, and

$$\text{rank}(\Theta(\sigma)) = \frac{1}{2}(n - v(\Psi(\sigma))) = n - \frac{1}{2}(n + v(\Psi(\sigma))) = n - \text{lis}(\sigma).$$

This implies the result. \square

4. Proof of Proposition 3

Let us first consider only fixed points in a permutation $\sigma \in \mathcal{S}_n$. Observe that if $\sigma \in \mathcal{S}_n(321)$ and $\sigma(i) = i$, then $(\sigma(1), \sigma(2), \dots, \sigma(i-1))$ is a permutation of $\{1, 2, \dots, i-1\}$, and $(\sigma(i+1), \sigma(i+2), \dots, \sigma(n))$ is a permutation of $\{i+1, i+2, \dots, n\}$. Indeed, if $\sigma(j) > i$ for some $j < i$, then necessarily $\sigma(k) < i$ for some $k > i$, and $(\sigma(j), \sigma(i), \sigma(k))$ would be an occurrence of 321.

Therefore, when we apply RSK to σ , the elements $\sigma(i), \sigma(i+1), \dots, \sigma(n)$ will never bump any of the elements $\sigma(1), \sigma(2), \dots, \sigma(i-1)$. In particular, the subtableaux of P and Q determined by the entries that are smaller than i will have the same shape. Furthermore, when the elements greater than i are placed in P and Q , the rows in which they are placed are independent of the subpermutation $(\sigma(1), \sigma(2), \dots, \sigma(i-1))$. Note also that $\sigma(i)$ will never be bumped.

When the Dyck path $\Psi(\sigma)$ is built from P and Q , this translates into the fact that the steps corresponding to $\sigma(i)$ in P and to i in Q will be, respectively, an up-step in the first half and a down-step in the second half, both at the same height and at the same distance from the center of the path. Besides, the part of the path between them will be itself the Dyck path corresponding to $(\sigma(i+1) - i, \sigma(i+2) - i, \dots, \sigma(n) - i)$. So, the fixed point $\sigma(i) = i$ determines a centered tunnel in $\Psi(\sigma)$. It is clear that the converse is also true, that is, every centered tunnel comes from a fixed point. This shows that $\text{fp}(\sigma) = \text{ct}(\Psi(\sigma))$, proving the first part of Proposition 3.

Let us now consider excedances in a permutation $\sigma \in \mathcal{S}_n(321)$. Our goal is to show that the excedances of σ correspond to right tunnels of $\Psi(\sigma)$. The first observation is that we can assume without loss of generality that σ has no fixed points. Indeed, if $\sigma(i) = i$ is a fixed point of σ , then the above reasoning shows that we can decompose $\Psi(\sigma) = AuBdC$, where AC is the Dyck path $\Psi((\sigma(1), \sigma(2), \dots, \sigma(i-1)))$ and B is a translation of the Dyck path $\Psi((\sigma(i+1) - i, \dots, \sigma(n) - i))$. But we have that $\text{exc}(\sigma) = \text{exc}((\sigma(1), \sigma(2), \dots, \sigma(i-1))) + \text{exc}((\sigma(i+1) - i, \dots, \sigma(n) - i))$ and $\text{rt}(AuBdC) = \text{rt}(AC) + \text{rt}(B)$, so in this case the result holds by induction on the number of fixed points. Note also that the above argument showed that $\text{fp}(\sigma) = \text{fp}((\sigma(1), \sigma(2), \dots, \sigma(i-1))) + \text{fp}((\sigma(i+1) - i, \dots, \sigma(n) - i)) + 1$ and $\text{ct}(AuBdC) = \text{ct}(AC) + \text{ct}(B) + 1$.

Suppose that $\sigma \in \mathcal{S}_n(321)$ has no fixed points. It is known that a permutation is 321-avoiding if and only if both the subsequence determined by its excedances and the one determined by the remaining elements (in this case, the antiexcedances) are increasing (see e.g. [8]). Denote by $X_i := (i, \sigma(i))$ the crosses of the array representation of σ . To simplify the presentation, we will refer indistinctively to i

or X_i , hoping this does not lead to a confusion. For example, we will say “ X_i is an excedance”, etc.

Define a matching between excedances and antiexcedances of σ by the following algorithm. Let $\sigma(i_1) < \sigma(i_2) < \dots < \sigma(i_k)$ be the excedances of σ and let $\sigma(j_1) < \sigma(j_2) < \dots < \sigma(j_{n-k})$ be the antiexcedances.

Matching Algorithm

- (1) Initialize $a := 1, b := 1$.
- (2) Repeat until $a > k$ or $b > n - k$:
 - (a) If $i_a > j_b$, then $b := b + 1$. (X_{j_b} is not matched.)
 - (b) Else if $\sigma(i_a) < \sigma(j_b)$, then $a := a + 1$. (X_{i_a} is not matched.)
 - (c) Else, match X_{i_a} with X_{j_b} ; $a := a + 1, b := b + 1$.
- (3) Output the matching sequence.

Example. Let $\sigma = (4, 1, 2, 5, 7, 8, 3, 6, 11, 9, 10)$ as in Fig. 5 below. We have $i_1 = 1, i_2 = 4, i_3 = 5, i_4 = 6, i_5 = 9$, and $j_1 = 2, j_2 = 3, j_3 = 7, j_4 = 8, j_5 = 10, j_6 = 11$. In the first execution of the loop in step (2) of the algorithm, neither $i_1 > j_1$ nor $\sigma(i_1) < \sigma(j_1)$ hold, so $X_{i_1} = (1, 4)$ and $X_{j_1} = (2, 1)$ are matched. Now we repeat the loop with $a = b = 2$, and since $i_2 > j_2$, we are in the case given by (2a) ($X_{j_2} = (3, 2)$ is not matched). In the next iteration, $a = 2$ and $b = 3$, so we match $X_{i_2} = (4, 5)$ and $X_{j_3} = (7, 3)$. Now we have $a = 3$ and $b = 4$, so we match $X_{i_3} = (5, 7)$ and $X_{j_4} = (8, 6)$. The values of a and b in the next iteration are 4 and 5, respectively, so we are in the case of (2b), $\sigma(i_4) = 8 < 9 = \sigma(j_5)$, and $X_{i_4} = (6, 8)$ is unmatched. Now $a = b = 5$, and we match $X_{i_5} = (9, 11)$ and $X_{j_6} = (10, 9)$. The matching algorithm ends here because now $a = 6 > 5 = k$.

An informal, more geometrical description of the matching algorithm is the following. For each pair of crosses of the array (seen as embedded in the plane), consider the line that they determine. If one of these lines has positive slope and leaves all the remaining crosses to the right, match the two crosses that determine it, and delete them from the array. If there is no line with these properties, delete the cross that is closer to the upper-left corner of the array (it is unmatched). Repeat the process until no crosses are left.




Fig. 5. Example of the matching for $\sigma = (4, 1, 2, 5, 7, 8, 3, 6, 11, 9, 10)$, and $\Psi(\sigma)$.

Now we consider the matched excedances on one hand and the unmatched ones on the other. We summarize rather technical results in the following two lemmas, which are proved in Section 5.

Lemma 5. *The following quantities are equal:*

- (1) *the number of matched pairs (X_i, X_j) , where X_i is an excedance and X_j an antiexcedance;*
- (2) *the length of the second row of P (or Q);*
- (3) *the number of right-side tunnels of $\Psi(\sigma)$;*
- (4) *the number of left-side tunnels of $\Psi(\sigma)$;*
- (5) $\frac{1}{2}(n - v(\Psi(\sigma)))$;
- (6) $n - \text{lis}(\sigma)$.

Note that (5) = (6) implies that $\text{lis}(\sigma) = \frac{1}{2}(n + v(\Psi(\sigma)))$, which is the third part of Proposition 3.

Lemma 6. *The number of unmatched excedances (resp. antiexcedances) of σ equals the number of right-across (resp. left-across) tunnels of $\Psi(\sigma)$.*

Since each excedance of σ either is part of a matched pair (X_i, X_j) or is unmatched, Lemmas 5 and 6 imply that the total number $\text{exc}(\sigma)$ of excedances equals the number of right-side tunnels of $\Psi(\sigma)$ plus the number of right-across tunnels, which is $\text{rt}(\Psi(\sigma))$. This implies the second part of Proposition 3.

To summarize, we have shown that the bijection Ψ satisfies all three properties described in the proposition. This completes the proof. \square

5. Proofs of the lemmas

Proof of Lemma 5. From the descriptions of the RSK algorithm and the matching, it follows that an excedance X_i and an antiexcedance X_j are matched with each other precisely when $\sigma(j)$ bumps $\sigma(i)$ when RSK is performed on σ , and that these are the only bumpings that take place. Indeed, an excedance never bumps anything because it is larger than the elements inserted before. On the other hand, when an antiexcedance X_j is inserted, it bumps the smallest element larger than $\sigma(j)$ which has not been bumped yet (which corresponds to an excedance that has not been matched yet), if such an element exists. This proves equality (1) = (2).

To see that (2) = (3), observe that right-side tunnels correspond to up-steps in the right half of $\Psi(\sigma)$, which by the construction of the bijection Ψ correspond to elements on the second row of Q . Equality (3) = (5) follows easily by counting the number of up-steps and down-steps of the right half of the path. Equality (4) = (5) is analogous.

Finally, Schensted's theorem states that the size of the first row of P equals the length of a longest increasing subsequence of σ (see [11] or [13, Section 7.23]). This implies that $(2)=(6)$, which completes the proof. \square

The reasoning used in the above proof gives a nice equivalent description of the recording tableau Q in terms of the array and the matching. Read the rows of the array from top to bottom. For i from 1 to n , place i on the first row of Q if X_i is an excedance or it is unmatched, and place i on the second row if X_i is a matched antiexcedance. In the construction of the right half of $\Psi(\sigma)$, this translates into drawing the path from right to left while reading the array from top to bottom, adjoining an up-step for each matched antiexcedance and a down-step for each other kind of cross.

To get a similar description of the tableau P , we use duality. By construction of the matching algorithm, the matching in the output is invariant under transposition of the array (reflection along the main diagonal). Recall the duality of the RSK correspondence: if $\text{RSK}(\sigma) = (P, Q)$, then $\text{RSK}(\sigma^{-1}) = (Q, P)$ (see e.g. [13, Section 7.13]). Therefore, the tableau P can be obtained by reading the columns of the array of σ from left to right and placing integers in P according to the following rule. For each column j , place j on the first row of P if the cross in column j is an antiexcedance or it is unmatched. Similarly, place j on the second row if the cross is a matched excedance. Equivalently, the left half of $\Psi(\sigma)$, from left to right, is obtained by reading the array from left to right and adjoining a down-step for each matched excedance, and an up-step for each of the remaining crosses.

In particular, when the left half of the path is constructed in this way, every matched pair (X_i, X_j) produces an up-step and a down-step, giving the latter a left-side tunnel. Similarly, in the construction of the right half of the path, a matched pair gives a right-side tunnel.

Proof of Lemma 6. It is enough to prove it only for the case of excedances. The case of antiexcedances follows from it considering σ^{-1} and noticing that the path $\Psi(\sigma^{-1})$ is obtained by reflecting $\Psi(\sigma)$ in a vertical axis through the middle of the path (this follows immediately from the duality of RSK). Let X_k be an unmatched excedance of σ . We use the above description of $\Psi(\sigma)$ in terms of the array and the matching. Each cross X_i produces a step r_i in the right half of the Dyck path and another step ℓ_i in the left half. Crosses above X_k produce steps to the right of r_k , and crosses to the left of X_k produce steps to the left of ℓ_k . In particular, there are $k - 1$ steps to the right of r_k , and $\sigma(k) - 1$ steps to the left of ℓ_k . Note that since X_k is an excedance and σ is 321-avoiding, all the crosses above it are also to the left of it. Consider the crosses that lie to the left of X_k . They can be of the following four kinds:

- *Unmatched excedances* X_i . They will necessarily lie above X_k , because the subsequence of excedances of σ is decreasing. Each one of these crosses contributes an up-step to the left of ℓ_k and down-step to the right of r_k .

- *Unmatched antiexcedances* X_j . They also have to lie above X_k , otherwise X_k would be matched with one of them. So, each such X_j contributes an up-step to the left of ℓ_k and down-step to the right of r_k .
- *Matched pairs* (X_i, X_j) (i.e. X_i is an excedance and X_j an antiexcedance), where both X_i and X_j lie above X_k . Both crosses together will contribute an up-step and a down-step to the left of ℓ_k , and an up-step and a down-step to the right of r_k .
- *Matched pairs* (X_i, X_j) (i.e. X_i is an excedance and X_j an antiexcedance), where X_j lies below X_k . The pair will contribute an up-step and a down-step to the left of ℓ_k . However, to the right of r_k , the only contribution will be a down-step produced by X_i .

Note that there cannot be an antiexcedance X_j to the left of X_k matched with an excedance to the right of X_k , because in this case X_j would have been matched with X_k by the algorithm. In the first three cases, the contribution to both sides of the Dyck path is the same, so that the heights of r_k and ℓ_k are equally affected. But since $\sigma(k) > k$, at least one of the crosses to the left of X_k must be below it, and this must be a matched antiexcedance as in the fourth case. This implies that the step r_k is at a higher y coordinate than ℓ_k . Let h_k be the height of ℓ_k . We now show that $\Psi(\sigma)$ has a right-across tunnel at height h_k .

Observe that h_k is the number of unmatched crosses to the left of X_k , and that the height of r_k is the number of unmatched crosses above X_k (which equals h_k) plus the number of excedances above X_k matched with antiexcedances below X_k . The part of the path between ℓ_k and the middle always remains at a height greater than h_k . This is because the only possible down-steps in this part can come from matched excedances X_i to the right of X_k , but then such a X_i is matched with an antiexcedance X_j to the right of X_k but to the left of X_i , which produces an up-step compensating the down-step associated to X_i . Similarly, the part of the path between r_k and the middle remains at a height greater than h_k . This is because the h_k down-steps to the right of r_k that come from unmatched crosses above X_k do not have a corresponding up-step in the part of the path between r_k and the middle. Hence, ℓ_k is the left end of a right-across tunnel, since the right end of this tunnel is to the right of r_k , which in turn is closer to the right end of $\Psi(\sigma)$ than ℓ_k is to its left end.

It can easily be checked that the converse is also true, namely that in every right-across tunnel of $\Psi(\sigma)$, the step at its left end corresponds to an unmatched excedance of σ . \square

6. Further applications

(6.1) Recall the result in [10] that the number of permutations $\sigma \in \mathcal{S}_n(132)$ (or $\sigma \in \mathcal{S}_n(321)$) with no fixed points is the *Fine number* F_n . This sequence is most easily defined by its relation to Catalan numbers:

$$C_n = 2F_n + F_{n-1} \quad \text{for } n \geq 2, \quad \text{and } F_1 = 0, F_2 = 1.$$

Although defined awhile ago, Fine numbers have received much attention in recent years (see a survey [3]). Special cases of our results give simple bijections between these two combinatorial interpretations of Fine numbers and a new one: the set of Dyck paths without centered tunnels. In particular, we obtain a bijective proof of the following result.

Corollary 7. *The number of Dyck paths $D \in \mathcal{D}_n$ without centered tunnels is equal to F_n .*

An analytical proof of this corollary can be easily deduced by combining results on Dyck paths in [4] with a combinatorial interpretation of Fine numbers given in [10]. However, ours is the first bijective proof of Corollary 7.

(6.2) We can also extend Propositions 3 and 4 to statistics $v_c(D)$ defined as the height at $x = n - c$ of the Dyck path $D \in \mathcal{D}_n$, for any $c \in \{0, \pm 1, \pm 2, \dots, \pm(n-1)\}$. The corresponding statistics in $\mathcal{S}_n(132)$ and in $\mathcal{S}_n(321)$ are generalizations of the rank of a permutation and the length of the longest increasing subsequence in a certain subpermutation of σ . The corresponding generalization of Theorem 2 is straightforward and is left to the reader.

(6.3) Let us also note that the limiting distribution of $\text{lis}(\cdot)$ on $\mathcal{S}_n(321)$ has been studied in [1]. From Theorem 2, the results in [1] can be translated into results on the limiting distribution of $\text{rank}(\cdot)$ on $\mathcal{S}_n(132)$.

(6.4) Our final application has appeared unexpectedly after the results of this paper have been obtained. We say that a permutation $\sigma \in \mathcal{S}_n$ is an *involution* if $\sigma = \sigma^{-1}$. In a recent paper [2] the authors introduce a notion of *refined restricted involutions* by considering the “number of fixed points” statistic on involutions avoiding different patterns $\pi \in \mathcal{S}_3$. They prove the following result:

Theorem 8 (Deutsch [2]). *The number of 321-avoiding involutions $\sigma \in \mathcal{S}_n$ with $\text{fp}(\sigma) = i$ equals the number of 132-avoiding involutions $\sigma \in \mathcal{S}_n$ with $\text{fp}(\sigma) = i$, for any $0 \leq i \leq n$.*

Let us show that Theorem 8 follows easily from our investigation. Indeed, for every Dyck path $D \in \mathcal{D}_n$ denote by D^* the path obtained by reflection of D from a vertical line $x = n$. Now observe that if $\Phi(\sigma) = D$, then $\Phi(\sigma^{-1}) = D^*$. Similarly, if $\Psi(\sigma) = D$, then $\Psi(\sigma^{-1}) = D^*$ (by the duality of RSK). Therefore, $\sigma \in \mathcal{S}_n(321)$ is an involution if and only if so is $\Theta(\sigma) \in \mathcal{S}_n(132)$, which implies the result. Furthermore, we obtain the following extension of Theorem 8.

Theorem 9. *The number of 321-avoiding involutions $\sigma \in \mathcal{S}_n$ with $\text{fp}(\sigma) = i$, $\text{exc}(\sigma) = j$ and $\text{lis}(\sigma) = k$ equals the number of 132-avoiding involutions $\sigma \in \mathcal{S}_n$ with $\text{fp}(\sigma) = i$, $\text{exc}(\sigma) = j$ and $\text{rank}(\sigma) = n - k$, for any $0 \leq i, j, k \leq n$.*

We leave the easy details of the proof to the reader.

Acknowledgments

We would like to thank Richard Stanley for suggesting the problem of enumerating excedances in pattern-avoiding permutations and for encouragement. We are also grateful to Emeric Deutsch, Peter McNamara and Aaron Robertson for helpful comments. S. Elizalde was partially supported by a MAE fellowship. I. Pak was supported by the NSA and the NSF.

References

- [1] E. Deutsch, A.J. Hildebrand, H.S. Wilf, Longest increasing subsequences in pattern-restricted permutations, *Electron J. Combin.* 9 (2) (2002,2003) #R12.
- [2] E. Deutsch, A. Robertson, D. Saracino, Refined restricted involutions, preprint, arxiv:math.CO/0212267.
- [3] E. Deutsch, L. Shapiro, A survey of the Fine numbers, *Discrete Math.* 241 (2001) 241–265.
- [4] S. Elizalde, Fixed points and excedances in restricted permutations, Proceedings of FPSAC Linköping University, Sweden, 2003.
- [5] M. Fulmek, Enumeration of permutations containing a prescribed number of occurrences of a pattern of length 3, *Adv. Appl. Math.* 30 (2003) 607–632.
- [6] D. Knuth, *The Art of Computer Programming*, Vol. III, Addison-Wesley, Reading, MA, 1973.
- [7] C. Krattenthaler, Permutations with restricted patterns and Dyck paths, *Adv. Appl. Math.* 27 (2001) 510–530.
- [8] A. Reifegerste, On the diagram of 132-avoiding permutations, preprint, arxiv:math.CO/0208006.
- [9] D. Richards, Ballot sequences and restricted permutations, *Ars Combin.* 25 (1988) 83–86.
- [10] A. Robertson, D. Saracino, D. Zeilberger, Refined restricted permutations, *Ann. Combin.* 6 (2003) 427–444.
- [11] C. Schensted, Longest increasing and decreasing subsequences, *Canad. J. Math.* 13 (1961) 179–191.
- [12] R. Simion, F.W. Schmidt, Restricted permutations, *European J. Combin.* 6 (1985) 383–406.
- [13] R. Stanley, *Enumerative Combinatorics*, Vols. I, II, Cambridge University Press, Cambridge, 1997, 1999.
- [14] J. West, Generating trees and the Catalan and Schröder numbers, *Discrete Math.* 146 (1995) 247–262.