CLASSIFICATION OF BIJECTIONS BETWEEN
321- AND 132-AVOIDING PERMUTATIONS

ANDERS CLAESSON AND SERGEY KITAEV

ABSTRACT. It is well-known, and was first established by Knuth in 1969, that the number of 321-avoiding permutations is equal to that of 132-avoiding permutations. In the literature one can find many subsequent bijective proofs of this fact. It turns out that some of the published bijections can easily be obtained from others. In this paper we describe all bijections we were able to find in the literature and show how they are related to each other via "trivial" bijections. We classify the bijections according to statistics preserved (from a fixed, but large, set of statistics), obtaining substantial extensions of known results. Thus, we give a comprehensive survey and a systematic analysis of these bijections.

We also give a recursive description of the algorithmic bijection given by Richards in 1988 (combined with a bijection by Knuth from 1969). This bijection is equivalent to the celebrated bijection of Simion and Schmidt (1985), as well as to the bijection given by Krattenthaler in 2001, and it respects 11 statistics—the largest number of statistics any of the bijections respects.

CONTENTS

1. INTRODUCTION AND MAIN RESULTS

Given two different bijections between two sets of combinatorial objects, what does it mean to say that one bijection is better than the other? Perhaps, a reasonable answer would be "The one that is easier to describe." While the ease of description and how easy it is to prove properties of the bijection using the description is one aspect to consider, an even more important aspect, in our opinion, is how well the bijection reflects and translates properties of elements of the respective sets.

1

A natural measure for a bijection between two sets of permutations, then, is how many statistics the bijection preserves. Obviously, we don't have an exhaustive list of permutation statistics, but we have used the following list as our "base" set:

$$\text{asc, des, exc, ldr, rdr, lir, rir, zeil, comp, lmax, lmin, rmax, rmin,}$$

$$\text{head, last, peak, valley, lds, lis, rank, cyc, fix, slmax} .$$

These are defined in Section 2. To make sure we find all statistics that a given bijection "essentially" preserves, we include in our list of statistics those that are obtained from our "basic" statistics by applying to them the *trivial bijections* on permutations (*reverse=r*, *complement=c*, *inverse=i*) and their compositions. Moreover, for each statistic stat, in this extended list we consider two other statistics: n-stat$(\pi) = n - \text{stat}(\pi)$ and m-stat$(\pi) = n + 1 - \text{stat}(\pi)$, where n is the length of the permutation. The meaning of n-stat or m-stat is often "non-stat"; for example, n-fix counts non-fixed-points.

This way each basic statistic gives rise to 24 statistics. The base set contains 23 statistics, giving a total of 552 statistics. There are, however, many statistics in that set that are equal as functions; for instance, des = asc.r, and peak = peak.r = valley.c, where we use a dot to denote composition of functions. Choosing one representative from each of the classes of equal statistics results in a final set of 190 statistics; we call this set STAT. In practice we settled for "empirical equality" when putting together STAT: we considered two statistics equal if they gave the same value on all 5914 permutations of length at most 7.

In the theorems below, the statistics presented are linearly independent. An example of linear dependence among the statistics over permutations avoiding 132 is lmin − lmax + n-des − head = 0. The results below are also maximal in that they cannot be non-trivially extended using statistics from STAT. That is, adding one more pair of equidistributed statistics from STAT to any of the results would create a linear dependency among the statistics.

A permutation $\pi = a_1 a_2 \ldots a_n$ avoids the *pattern* 321 if there are no indices $i < j < k$ such that $a_k < a_j < a_i$. It avoids 132 if there are no indices $i < j < k$ such that $a_i < a_k < a_j$. Avoidance of other patterns is defined similarly.

Knuth [6, 7] showed that the number of permutations avoiding a pattern of length 3 is independent of the pattern. This number is the *n-th Catalan number*, $C_n = \frac{1}{n+1}\binom{2n}{n}$. To prove this it suffices, due to the symmetry afforded by the trivial bijections on permutations, to consider one representative from $\{123, 321\}$ and one from $\{132, 231, 213, 312\}$. That symmetry also means that to prove this bijectively, it suffices to find a bijection from the set of permutations avoiding a pattern in one of the classes to permutations avoiding a pattern in the other. This turns out to be a rather complicated problem. Several authors have, however, succeeded in doing so [4, 5, 8, 10, 11, 13, 14, 15]; we call those bijections

Knuth, Knuth-Rotem, Simion-Schmidt, Knuth-Richards, West,

Krattenthaler, Reifegerste, Elizalde-Deutsch, and Mansour-Deng-Du .

They are described in Section 3. In Section 4 we define, using recursion, a "new" bijection called Φ. It turns out that Φ is related, via trivial bijections, to the bijection by Knuth and Richards.

The main results of this paper are contained in the following three theorems. The first theorem substantially extends what was previously known about statistics

preserved by the bijections. In bold we mark the results known before this paper—there is a total of 7 pairs of those; the remaining 68 pairs are new.

Theorem 1. *The following results are maximal in the sense that adding one more pair of equidistributed statistics from* STAT *to any of the results would create a linear dependency among the statistics. In bold we mark the results that were known; also, we indicate the sets between which a bijection acts.*

(11) **Knuth-Richards**, *132-avoiding permutations → 123-avoiding permutations*

valley.i	valley	lmin	ldr.i	head.i	comp.r	rank	ldr	lir.i	lir	rmax
valley	valley.i	lmin	ldr	head	comp.r	rank	ldr.i	slmax.c	slmax.i.r	head.i.r

(11) **Simion-Schmidt**, *123-avoiding permutations → 132-avoiding permutations*

valley	valley.i	**lmin**	ldr	head	comp.r	rank	ldr.i	slmax.c	slmax.i.r	head.i.r
valley	valley.i	**lmin**	ldr	head	comp.r	rank	ldr.i	lir	lir.i	rmin

(11) **Krattenthaler**, *123-avoiding permutations → 132-avoiding permutations*

peak.i	peak	rmax	zeil	last.i.r	comp.r	rank.r.c	rdr	slmax.r.i	slmax.r	last
valley	valley.i	lmin	ldr	head	comp.r	rank	ldr.i	lir	lir.i	rmin

(11) **Mansour-Deng-Du**, *321-avoiding permutations → 231-avoiding permutations*

valley	peak.i	rmin	rir	last	comp	rank.r	lir.i	slmax.c.r	slmax.i	head.i
valley	peak.i	rmin	rir	last	comp	rank.r	lir.i	rdr	ldr.i	lmin

(9) **Knuth-Rotem**, *321-avoiding permutations → 132-avoiding permutations*

valley.i	peak	**exc**	slmax	head	slmax.r.c.i	rir.i	lir	last.i
valley.i	valley	**des**	rdr	ldr.i	zeil	lmax	rmin	m-ldr

(9) **Reifegerste**, *321-avoiding permutations → 132-avoiding permutations*

valley	peak.i	**exc**	slmax.i	head.i	slmax.r.c	rir	lir.i	last
valley	valley.i	**des**	zeil	ldr	rdr	rmin	lmax	m-ldr.i

(7) **West**, *123-avoiding permutations → 132-avoiding permutations*

valley.i	exc.r	slmax.i.r	slmax.c	ldr	ldr.i	head
valley.i	asc	lir.i	comp	rmax	ldr.i	head

(5) **Knuth**, *321-avoiding permutations → 132-avoiding permutations*

exc	**fix**	lir.i	lir	**lis**
exc	**fix**	rmin	lmax	**n-rank**

(1) **Elizalde-Deutsch**, *321-avoiding permutations → 132-avoiding permutations*

fix
fix

The numbers in parenthesis in Theorem 1 indicate the number of statistics respected. It turns out that bijections with the same number are "trivially" related. The next theorem makes this precise.

Theorem 2. *The following relations among bijections between* 321- *and* 132-*avoiding permutations hold:*

$$\text{reverse} \circ \Phi^{-1} \; = \; \text{inverse} \circ \text{Simion-Schmidt} \circ \text{reverse}$$
$$= \; \text{inverse} \circ \text{Krattenthaler} \circ \text{reverse} \circ \text{inverse}$$
$$= \; \text{inverse} \circ \text{reverse} \circ \text{Mansour-Deng-Du}$$
$$= \; \text{Knuth-Richards}^{-1} \circ \text{reverse}$$

and

$$\text{Reifegerste} \; = \; \text{inverse} \circ \text{Knuth-Rotem} \circ \text{inverse}$$

Also, there are no other relations among the bijections and their inverses via the trivial bijections that does not follow from the ones above.

Thus if we regard all bijections as bijections from 321- to 132-avoiding permutations—applying the transformations in Theorem 2—then we get the following condensed version of Theorem 1.

Theorem 3. *For bijections from 321- to 132-avoiding permutations we have the following equidistribution results. These results are maximal in the sense that adding one more pair of equidistributed statistics from* STAT *to any of the results would create a linear dependency among the statistics.*

(11) Φ, Knuth-Richards, Krattenthaler, Mansour-Deng-Du, Simion-Schmidt

valley	peak.i	rmin	rir	last	comp	rank.r	lir.i	slmax.c.r	slmax.i	head.i
valley	valley.i	lmin	ldr	head	comp.r	rank	ldr.i	lir	lir.i	rmin

(9) Knuth-Rotem, Reifegerste

valley	peak.i	exc	slmax.i	head.i	slmax.r.c	rir	lir.i	last
valley	valley.i	des	zeil	ldr	rdr	rmin	lmax	m-ldr.i

(7) West

peak.i	exc	slmax.i	slmax.r.c	rir	lir.i	last
valley.i	asc	lir.i	comp	rmax	ldr.i	head

(5) Knuth

exc	fix	lir.i	lir	lis
exc	fix	rmin	lmax	n-rank

(1) Elizalde-Deutsch

fix
fix

In Section 2 we define the relevant statistics; in Section 3 we describe the bijections; in Section 4 we give a new recursive description of the bijection by Knuth and Richards; in Section 5 we prove Theorem 2; and in Section 6 we prove Theorem 1.

2. Permutation statistics

The permutation π on $\{1, 2, \ldots, n\}$ that sends 1 to a_1, 2 to a_2, etc, we denote $\pi = a_1 a_2 \ldots a_n$, and we call a_i the i-th letter of π. A permutation statistic is simply a function from permutations to $\mathbb{N}$. For example, the permutation statistic asc is defined thus: An *ascent* in π is a letter that is followed by a larger letter; in other words, an a_i such that $a_i < a_{i+1}$. By $\text{asc}(\pi)$ we denote the number of ascents in π.

Similarly, a *descent* is a letter followed by a smaller letter, and by $\mathrm{des}(\pi)$ we denote the number of descents in π.

For words α and β over the alphabet $\mathbb{N}$ we define that $\alpha \prec \beta$ if for all letters a in α and all letters b in β we have $a < b$. For instance, $42 \prec 569$. A *component* of π is a nonempty segment τ of π such that $\pi = \sigma\tau\rho$ with $\sigma \prec \tau \prec \rho$, and such that if $\tau = \alpha\beta$ and $\alpha \prec \beta$ then α or β is empty. By $\mathrm{comp}(\pi)$ we denote the number of components of π. For instance, $\mathrm{comp}(213645) = 3$, the components being 21, 3, and 645.

A *left-to-right minimum* of π is a letter with no smaller letter to the left of it; the number of left-to-right minima is denoted $\mathrm{lmin}(\pi)$. The statistics *right-to-left minima* (rmin), *left-to-right maxima* (lmax), and *right-to-left maxima* (rmax) are defined similarly.

In the following table we define the remaining statistics that are of interest to us. For reference we include the statistics already defined in the preceding few paragraphs.

$$
\begin{aligned}
\mathrm{asc} &= \text{number of ascent;} \\
\mathrm{comp} &= \text{number of components;} \\
\mathrm{des} &= \text{number of descents;} \\
\mathrm{exc} &= \text{number of excedances: positions } i \text{ such that } a_i > i; \\
\mathrm{fix} &= \text{number of fixed points: positions } i \text{ such that } a_i = i; \\
\mathrm{head} &= \text{first element: head}(\pi) = a_1; \\
\mathrm{last} &= \text{last element: last}(\pi) = a_n;
\end{aligned}
$$

ldr = length of the leftmost decreasing run: largest i such that $a_1 > a_2 > \cdots > a_i$;

lds = length of the longest decreasing sequence in a permutation;

lir = length of the leftmost increasing run: largest i such that $a_1 < a_2 < \cdots < a_i$;

lis = length of the longest increasing sequence in a permutation;

lmax = number of left-to-right maxima;

lmin = number of left-to-right minima;

peak = number of peaks: positions i in π such that $a_{i-1} < a_i > a_{i+1}$;

rank = largest k such that $a_i > k$ for all $i \leq k$ (see [5]);

rdr = lir.r = length of the rightmost decreasing run;

rmax = number of right-to-left maxima;

rmin = number of right-to-left minima;

rir = ldr.r = length of the rightmost increasing run;

slmax = the number of letters to the left of second left-to-right maximum in $\pi\infty$: largest i such that $a_1 \geq a_1, a_1 \geq a_2, \ldots, a_1 \geq a_i$;

valley = number of valleys: positions i in π such that $a_{i-1} > a_i < a_{i+1}$;

zeil = rdr.i = length of the longest subword $n(n-1)\ldots i$ (see [16]).

Let us also describe some of the derived statistics:

comp.r = number of reverse components: a *reverse component* is a nonempty segment τ of π such that $\pi = \sigma\tau\rho$ with $\sigma \succ \tau \succ \rho$, and such that if $\tau = \alpha\beta$ and $\alpha \succ \beta$ then α or β is empty;

head.i $=$ position of the smallest letter;

last.i $=$ position of the largest letter;

lir.i $=$ zeil.c $=$ largest i such that $12\ldots i$ is a subword in π;

peak.i $=$ number of letters a_i that are to the right of both $a_i - 1$ and $a_i + 1$;

valley.i $=$ number of letters a_i that are to the left of both $a_i - 1$ and $a_i + 1$.

3. Bijections in the literature

In this section we describe the bijections, and we try to stay close to the original sources when doing so. In what follows $\mathcal{S}_n(\tau)$ is the set of τ-avoiding permutations of length n, and $\mathcal{D}_n$ is the set of Dyck paths of length $2n$.

3.1. **Knuth's bijection, 1973.**

Knuth [6, pp. 242–243] gives a bijection from 312-avoiding permutations to "stack words". Formulated a bit differently, it amounts to a bijection from 132-avoiding permutations to Dyck paths. Knuth [7, pp. 60–61] also gives a bijection from 321-avoiding permutations to Dyck paths. By letting permutations that are mapped to the same Dyck path correspond to each other, a bijection between 321- and 132-avoiding permutation is obtained—we call it Knuth's bijection.

We start by describing the bijection from 132-avoiding permutations to Dyck paths. We shall refer to it as the *standard bijection*. (This bijection is the same as the one given by Krattenthaler [8], who, however, gives a non-recursive description of it; see Section 3.6.) Let $\pi = \pi_L n \pi_R$ be a 132-avoiding permutation of length n. Each letter of π_L is larger than any letter of π_R, or else a 132 pattern would be formed. We define the standard bijection f recursively by $f(\pi) = uf(\pi_L)df(\pi_R)$ and $f(\epsilon) = \epsilon$. Here, and elsewhere, ϵ denotes the empty word/permutation. Thus, under the standard bijection, the position of the largest letter in a 132-avoiding permutation determines the first return to x-axis and vice versa. For instance,

$$f(7564213) = udf(564213) = uduf(5)df(4213) = uduuddudf(213)$$

$$= uduudduduf(21)d = uduudduduudf(1)d$$

$$= uduudduduududd$$

$$= \text{(image)}$$

As mentioned, Knuth also gives a bijection from 321-avoiding permutations to Dyck paths: Given a 321-avoiding permutation, start by applying the *Robinson-Schensted-Knuth correspondence* to it. This classic correspondence gives a bijection between permutations π of length n and pairs (P, Q) of *standard Young tableaux* of the same shape $\lambda \vdash n$. As is well known, the length of the longest decreasing subword in π corresponds to the number of rows in P (or Q). Thus, for 321-avoiding permutations, the tableaux P and Q have at most two rows.

The *insertion tableau* P is obtained by reading $\pi = a_1 a_2 \ldots a_n$ from left to right and, at each step, inserting a_i to the partial tableau obtained thus far. Assume that $a_1, a_2, \ldots, a_{i-1}$ have been inserted. If a_i is larger than all the elements in the first row of the current tableau, place a_i at the end of the first row. Otherwise, let m be the leftmost element in the first row that is larger than a_i. Place a_i in the square that is occupied by m, and place m at the end of the second row. The *recording tableau* Q has the same shape as P and is obtained by placing i, for i from 1 to n, in the position of the square that in the construction of P was created at step

i (when a_i was inserted). For example, the pair of tableaux corresponding to the 321-avoiding permutation 3156247 we get by the following sequence of insertions:

$$(\epsilon \mid \epsilon) \rightarrow (3 \mid 1) \rightarrow \begin{pmatrix} 1 & \Big| & 1 \\ 3 & \Big| & 2 \end{pmatrix}$$

$$\rightarrow \begin{pmatrix} 15 & \Big| & 13 \\ 3 & \Big| & 2 \end{pmatrix} \rightarrow \begin{pmatrix} 156 & \Big| & 134 \\ 3 & \Big| & 2 \end{pmatrix}$$

$$\rightarrow \begin{pmatrix} 126 & \Big| & 134 \\ 35 & \Big| & 25 \end{pmatrix} \rightarrow \begin{pmatrix} 124 & \Big| & 134 \\ 356 & \Big| & 256 \end{pmatrix} \rightarrow \begin{pmatrix} 1247 & \Big| & 1347 \\ 356 & \Big| & 256 \end{pmatrix}.$$

The pair of tableaux (P, Q) is then turned into a Dyck path D. The first half, A, of the Dyck path we get by recording, for i from 1 to n, an up-step if i is in the first row of P, and a down-step if it is in the second row. Let B be the word obtained from Q in the same way but interchanging the roles of u and d. Then $D = AB^r$ where B^r is the reverse of B. Continuing with the example above we get

$$D = \text{}$$

Elizalde and Pak [5] use this bijection together with a slight modification of the standard bijection to give a combinatorial proof of a generalization of the result by Robertson et al. [12] that fixed points have the same distribution on 123- and 132-avoiding permutations. The modification they use is to reflect the Dyck path obtained from the standard bijection with respect to the vertical line crossing the path in the middle. Alternatively, the path can be read from the permutation diagram as described in [5]. We follow Elizalde and Pak and apply the same modification. After reflection, the path $f(7564213)$ above is the same as the path D in the preceding example. Thus the image of the 321-avoiding permutation 3156247 under what we call Knuth's bijection is the 132-avoiding permutation 7564213.

3.2. **Knuth-Rotem's bijection, 1975.** Rotem [13] gives a bijection between 321-avoiding permutations and Dyck paths, described below. Combining it with the standard bijection gives a bijection from 321- to 132-avoiding permutations—we call it *Knuth-Rotem's bijection*.

A *ballot-sequence* $b_1 b_2 \ldots b_n$ satisfies the two conditions

(1) $b_1 \leq b_2 \leq \cdots \leq b_n$;
(2) $0 \leq b_i \leq i - 1$, for $i = 1, 2, \ldots, n$.

Let $\pi = p_1 p_2 \ldots p_n$ be a 321-avoiding permutation. From it we construct a ballot-sequences $b_1 b_2 \ldots b_n$: Let $b_1 = 0$. For $i = 2, \ldots, n$, let $b_i = b_{i-1}$ if p_i is a left-to-right maximum in π, and let $b_i = p_i$ otherwise.

For the permutation $\pi = 2513476$ we get the ballot-sequences 0013446. This sequence we represent by a "bar-diagram", which in turn can be viewed as a lattice

path from $(0,0)$ to $(7,7)$:

Rotating that path counter clockwise by $3\pi/4$ radians we get

In the previous subsection we saw that this path is $f(7564213)$ where f is the standard bijection. Thus the image of the 321-avoiding permutation 2513476 under Knuth-Rotem's bijection is the 132-avoiding permutation 7564213.

3.3. Simion-Schmidt's bijection, 1985.

Consider the following algorithm:

Input: A permutation $\sigma = a_1 a_2 \ldots, a_n$ in $\mathcal{S}_n(123)$.
Output: A permutation $\tau = c_1 c_2 \ldots c_n$ in $\mathcal{S}_n(132)$.

```
1        c₁ := a₁; x := a₁
2        for i = 2,…,n:
3            if aᵢ < x:
4                cᵢ := aᵢ; x := aᵢ
5            else:
6                cᵢ := min{ k | x < k ≤ n, k ≠ cⱼ for all j < i }
```

The map $\sigma \mapsto \tau$ is the Simion-Schmidt bijection [14]. As an example, the 123-avoiding permutation 6743152 maps to the 132-avoiding permutation 6743125.

3.4. Knuth-Richards' bijection, 1988.

Richards' bijection [11] from Dyck paths to 123-avoiding permutations is given by the following algorithm:

Input: A Dyck path $P = b_1 b_2 \ldots b_{2n}$.
Output: A permutation $\pi = a_1 a_2 \ldots a_n$ in $\mathcal{S}_n(123)$.

```
1        r := n + 1; s := n + 1; j := 1
2        for i = 1,…,n:
3            if bⱼ is an up-step:
4                repeat s := s − 1; j := j + 1 until bⱼ is a down-step
5                aₛ := i
6            else:
7                repeat r := r − 1 until aᵣ is unset
8                aᵣ := i
9            j := j + 1
```

The Knuth-Richards bijection, from $\mathcal{S}_n(132)$ to $\mathcal{S}_n(123)$, is defined by

$$\text{Knuth-Richards} \;=\; \text{Richards} \circ f,$$

where f is the standard bijection from 132-avoiding permutations to Dyck paths, and Richards is the algorithm just described. As an example, applying Knuth-Richards' bijection to 6743125 yields 5743612.

3.5. **West's bijection, 1995.** West's bijection [15] is induced by an isomorphism between *generating trees*. The two isomorphic trees generate 123- and 132-avoiding permutations, respectively. We give a brief description of that bijection: Given a permutation $\pi = p_1 p_2 \ldots p_{n-1}$ and a positive integer $i \le n$, let

$$\pi^i \;=\; p_1 \ldots p_{i-1} \, n \, p_i \ldots p_{n-1};$$

we call this *inserting n into site i*. With respect to a fixed pattern τ we call site i of π in $\mathcal{S}_{n-1}(\tau)$ *active* if the insertion of n into site i creates a permutation in $\mathcal{S}_n(\tau)$.

For $i = 0, \ldots, n-1$, let a_{i+1} be the number of active sites in the permutation obtained from π by removing the i largest letters. The *signature* of π is the word

$$a_0 a_1 \ldots a_{n-1}.$$

West [15] showed that for 123-avoiding permutations, as well as for 132-avoiding permutations, the signature determines the permutation uniquely. This induces a natural bijection between the two sets. For example, the 123-avoiding permutation 536142 corresponds to the 132-avoiding permutation 534612—both have the same signature, 343322.

3.6. **Krattenthaler's bijection, 2001.** Krattenthaler's bijection [8] uses Dyck paths as intermediate objects. Permutations that are mapped to the same Dyck path correspond to each other under this bijection.

The first part of Krattenthaler's bijection is a bijection from 123-avoiding permutations to Dyck paths. Reading right to left, let the right-to-left maxima in π be $m_1, m_2, \ldots, m_s$, so that

$$\pi \;=\; w_s m_s \ldots w_2 m_2 w_1 m_1,$$

where w_i is the subword of π in between m_{i+1} and m_i. Since π is 123-avoiding, the letters in w_i are in decreasing order. Moreover, all letters of w_i are smaller than those of w_{i+1}.

To define the bijection, read π from right to left. Any right-to-left maximum m_i is translated into $m_i - m_{i-1}$ up-steps (with the convention $m_0 = 0$). Any subword w_i is translated into $|w_i| + 1$ down-steps, where $|w_i|$ denotes the number of letters of w_i. Finally, the resulting path is reflected in a vertical line through the center of the path. Alternatively, we could have generated the Dyck path from right to left.

The second part of Krattenthaler's bijection is a bijection from 132-avoiding permutations to Dyck paths. Read $\pi = p_1 p_2 \ldots p_n$ in $\mathcal{S}_n(132)$ from left to right and generate a Dyck path. When p_j is read, adjoin, to the path obtained thus far, as many up-steps as necessary to reach height $h_j + 1$, followed by a down-step to height h_j (measured from the x-axis); here h_j is the number of letters in $p_{j+1} \ldots p_n$ which are larger than p_j. This procedure can be shown to be equivalent to the standard bijection from 132-avoiding permutations to Dyck paths.

For instance, Krattenthaler's bijection sends the permutation 536142 in $\mathcal{S}_6(123)$ to the permutation 452316 in $\mathcal{S}_6(132)$—both map to the same Dyck path,

3.7. **Reifegerste's bijection, 2002.**

	1	2	3	4	5	6	7	8
1							∘	
2						□		∘
3					∘			
4				□		∘		
5			□		∘			
6	□	∘						
7	∘							
8			∘					

This figure illustrates Reifegerste's bijection [10]. It pictures the 321-avoiding permutation $\pi = 13256847$ and the 132-avoiding permutation $\pi' = 78564213$, two permutations that correspond to each other under that bijection.

Let $\pi = a_1 a_2 \ldots a_n$ be a 321-avoiding permutation, and let E be the set of pairs

$$E = \{\, (i, a_i) \mid i \text{ is an excedance} \,\}.$$

For each pair (i, a_i) in E, we place a square, called an *E-square*, in position $(i, n + 1 - a_i)$ in an $n \times n$ *permutation matrix*. (E uniquely determines π.) Next we shade each square (a, b) of the matrix where there are no E-squares in the region $\{(i, j) \mid i \geq a, \ j \geq b\}$, thus obtaining a *Ferrer's diagram*. Finally, we get the 132-avoiding permutation π' corresponding to π by placing dots (circles), row by row starting from the first row, in the leftmost available shaded square such that there are no two dots in any column or row. If (i, j) contains a dot, then $\pi'(i) = j$.

3.8. **Elizalde-Deutsch's bijection, 2003.**
Here is an outline of a bijection by Elizalde and Deutsch [4]: Map 321- and 132-avoiding permutation bijectively to Dyck paths; use an automorphism Ψ on Dyck paths; and match permutations with equal paths.

We start by describing the automorphism Ψ. Let P be a Dyck path of length $2n$. Each up-step of P has a corresponding down-step in the sense that the path between the up-step and the down-step form a proper Dyck path. Match such pairs of steps. Let σ in $\mathcal{S}_{2n}$ be the permutation defined by $\sigma_i = (i + 1)/2$ if i is odd, and $\sigma_i = 2n + 1 - i/2$ otherwise. For i from 1 to $2n$, consider the σ_i-th step of P. If the corresponding matching step has not yet been read, define the i-th step of $\Psi(P)$ to be an up-step, otherwise let it be a down-step. For example,

$$\Psi(uuduudududdud) = uuudddduduuddud.$$

The bijection ψ from 321-avoiding permutations to $\mathcal{D}_n$ is defined as follows. Any permutation π in $\mathcal{S}_n$ can be represented as an $n \times n$ array with crosses in the squares $(i, \pi(i))$. Given the array of π in $\mathcal{S}_n(321)$, consider the path with *down-* and *right*-steps along the edges of the squares that goes from the upper-left corner to the lower-right corner of the array leaving all the crosses to the right and remaining always as close to the main diagonal as possible. Then the corresponding Dyck path is obtained from this path by reading an up-step every time the path moves down, and a down-step every time the path moves to the right. For example,

$$\psi(2314657) = uuduudududdud.$$

The bijection ϕ from 132-avoiding permutations to $\mathcal{D}_n$ is the standard bijection followed by a reflection of the path with respect to a vertical line through the middle of the path. For example,

$$\phi(7432516) = uuduudududdud.$$

The Elizalde-Deutsch bijection, from $\mathcal{S}_n(321)$ to $\mathcal{S}_n(132)$, is defined by

$$\text{Elizalde-Deutsch} = \phi^{-1} \circ \Psi^{-1} \circ \psi.$$

FIGURE 1.

As an example, it send 2314657 to 2314657.

3.9. **Mansour-Deng-Du's bijection, 2006.** Let i be a positive integer smaller than n. Let $s_i : \mathcal{S}_n \to \mathcal{S}_n$ act on permutations by interchanging the letters in positions i and $i+1$. We call s_i a *simple transposition*, and write the action of s_i as πs_i. So, $\pi(s_i s_j) = (\pi s_i)s_j$. For any permutation π of length n, the *canonical reduced decomposition* of π is

$$\pi = (12 \ldots n)\sigma = (12 \ldots n)\sigma_1 \sigma_2 \ldots \sigma_k,$$

where $\sigma_i = s_{h_i} s_{h_i - 1} \ldots s_{t_i}$, $h_i \geq t_i$, $1 \leq i \leq k$ and $1 \leq h_1 < h_2 < \cdots < h_k \leq n - 1$. For example, $415263 = (s_3 s_2 s_1)(s_4 s_3)(s_5)$.

Mansour, Deng and Du [9] use canonical reduced decompositions to construct a bijection between $\mathcal{S}_n(321)$ and $\mathcal{S}_n(231)$. They show that a permutation is 321-avoiding precisely when $t_i \geq t_{i-1} + 1$ for $2 \leq i \leq k$ [9, Thm. 2]. They also show that a permutation is 231-avoiding precisely when $t_i \geq t_{i-1}$ or $t_i \geq h_{i-j} + 2$ for $2 \leq i \leq k$ and $1 \leq j \leq i - 1$ [9, Thm. 15]. Using these two theorems they build their bijection, which is composed of two bijections: one from $\mathcal{S}_n(321)$ to $\mathcal{D}_n$, and one from $\mathcal{S}_n(231)$ to $\mathcal{D}_n$.

For a Dyck path P, we define the $(x + y)$-labeling of P as follows: each cell in the region enclosed by P and the x-axis, whose corner points are (i, j), $(i+1, j-1)$, $(i+2, j)$ and $(i+1, j+1)$ is labeled by $(i+j)/2$. If $(i-1, j-1)$ and (i, j) are starting points of two consecutive up-steps, then we call the cell with leftmost corner (i, j) an *essential cell* and the up-step $((i - 1, j - 1), (i, j))$ its *left arm*. We define the *zigzag strip* of P as follows: If there is no essential cell in P, then the zigzag strip is simply the empty set. Otherwise, we define the zigzag strip of P as the border strip that begins at the rightmost essential cell. For example, the zigzag strip of the Dyck path *uuduuududdddudduduuddud* in Figure 1 is the shaded cell labeled by 9, while for the Dyck path *uuduuududddd* (obtained from that in Figure 1 by ignoring the steps 15 to 22) the zigzag strip is the shaded connected cells labeled by 2, 3, 4, 5 and 6.

Let $P_{n,k}$ be a Dyck path of semi-length n containing k essential cells. We define its *zigzag decomposition* as follows: The zigzag decomposition of $P_{n,0}$ is the empty set. The zigzag decomposition of $P_{n,1}$ is the zigzag strip. If $k \geq 2$, then we decompose $P_{n,k} = P_{n,k-1}Q$, where Q is the zigzag strip of $P_{n,k}$ and $P_{n,k-1}$ is the Dyck path obtained from P by deleting Q. Reading the labels of Q from left to right, ignoring repetitions, we get a sequence of numbers $\{i, i + 1, \ldots, j\}$, and we associate Q with the sequence of simple decompositions $\sigma_k = s_j s_{j-1} \ldots s_i$. For $P_{n,i}$ with $i \leq k - 1$ repeat the above procedure to get σ_{k-1}, ..., σ_2, σ_1. The zigzag decomposition of $P_{n,k}$ is then given by $\sigma = \sigma_1 \sigma_2 \ldots \sigma_k$.

FIGURE 2.

From the zigzag decomposition we get a 321-avoiding permutation $\pi = (12 \ldots n)\sigma$ whose canonical reduced decomposition is σ. For the Dyck path $P_{11,4}$ in Figure 1 we have

$$\sigma = (s_3 s_2 s_1)(s_4 s_3)(s_6 s_5 s_4)(s_9)$$

and the corresponding permutation in $\mathcal{S}_{11}(321)$ is $(4, 1, 5, 7, 2, 3, 6, 8, 10, 9, 11)$.

We will now describe a map from Dyck paths to 231-avoiding permutations. For a Dyck path P, we define $(x - y)$-labeling of P as follows (this labeling seems to be considered for the first time in [1]): each cell in the region enclosed by P and the x-axis, whose corner points are (i, j), $(i + 1, j - 1)$, $(i + 2, j)$ and $(i + 1, j + 1)$ is labeled by $(i - j + 2)/2$. We define the *trapezoidal strip* of P as follows: If there is no essential cell in P, then the trapezoidal strip is simply the empty set. Otherwise, we define the trapezoidal strip of P as the horizontal strip that touches the x-axis and starts at the rightmost essential cell. For example, the trapezoidal strip of the Dyck path *uuduuududdddudduduuddud* in Figure 2 is the shaded cell labeled by 9, while for the Dyck path *uuduuududddd* (obtained from that in Figure 2 by ignoring the steps 15 to 22) the zigzag strip is the down-most shaded strip with labels 1, 2, 3, 4, 5 and 6.

Let $P_{n,k}$ be a Dyck path of semi-length n containing k essential cells. We define its *trapezoidal decomposition* as follows: The trapezoidal decomposition of $P_{n,0}$ is the empty set. The trapezoidal decomposition of $P_{n,1}$ is the trapezoidal strip. If $k \geq 2$, then we decompose $P_{n,k}$ into $P_{n,k} = Q_1 u Q_2 d$, where u is the left arm of the rightmost essential cell that touches the x-axis, d is the last down step of $P_{n,k}$, and Q_1 and Q_2 carry the labels in $P_{n,k}$. Reading the labels of the trapezoidal strip of $P_{n,k}$ from left to right we get a sequence $\{i, i+1, \ldots, j\}$, and we set $\sigma_k = s_j s_{j-1} \ldots s_i$. Repeat the above procedure for Q_1 and Q_2. Suppose the trapezoidal decomposition of Q_1 and Q_2 are σ' and σ'' respectively, then the trapezoidal decomposition for $P_{n,k}$ is $\sigma = \sigma' \sigma'' \sigma_k$.

From the trapezoidal decomposition we get a 231-avoiding permutation $\pi = (12 \ldots n)\sigma$ whose canonical reduced decomposition is σ. For the Dyck path $P_{11,4}$ in Figure 2 we have

$$\sigma = (s_3 s_2)(s_4 s_3 s_2)(s_6 s_5 s_4 s_3 s_2 s_1)(s_9)$$

and the corresponding permutation in $\mathcal{S}_{11}(231)$ is $(7, 1, 5, 4, 2, 3, 6, 8, 10, 9, 11)$.

The two maps involving Dyck paths described in this subsection induce a bijection from 321-avoiding to 231-avoiding permutations.

4. A RECURSIVE DESCRIPTION OF THE KNUTH-RICHARDS BIJECTION

We call a permutation π *indecomposable* if $\text{comp}(\pi) = 1$; otherwise we call π *decomposable*. Equivalently, if we define the sum $\oplus$ on permutations by $\sigma \oplus \tau = \sigma\tau'$,

where τ' is obtained from τ by adding $|\sigma|$ to each of its letters, then a permutation is indecomposable if it cannot be written as the sum of two nonempty permutations.

We shall describe, separately for 231- and 321-avoiding permutations, how to generate the indecomposable permutations, thus inducing a bijection we call Φ.

For a permutation of length n to be 231-avoiding everything to the left of n has to be smaller than anything to the right of n. Clearly, if there is at least one letter to the left of n, then the permutation is decomposable (everything to the right of n, including n, would form the last component). Thus a 231-avoiding permutation of length n is indecomposable if and only if it starts with n.

To build an indecomposable 231-avoiding permutation of length n from a 231-avoiding permutation of length $n-1$ we simply prepend n. Let us call this map α. For instance, $\alpha(2134) = 52134$.

Given k indecomposable 231-avoiding permutations $\pi_1, \pi_2, \ldots, \pi_k$, we build the corresponding permutation by summing: $\pi_1 \oplus \pi_2 \oplus \cdots \oplus \pi_k$. Given k indecomposable 321-avoiding permutations $\pi_1, \pi_2, \ldots, \pi_k$ we build the corresponding permutation by summing in reverse order: $\pi_k \oplus \pi_{k-1} \oplus \cdots \oplus \pi_1$.

Here is how we build an indecomposable 321-avoiding permutation π' of length n from a 321-avoiding permutation π of length $n-1$:

$$
\begin{array}{ccccccccccc}
\pi & = & & 2 & 4 & 1 & 3 & 5 & \boxed{7} & 6 & \boxed{9} & 8 \\
& & 2 & 4 & \boxed{10} & 1 & 3 & 5 & \boxed{7} & 6 & \boxed{9} & 8 \\
\pi' & = & 2 & 4 & \boxed{7} & 1 & 3 & 5 & \boxed{9} & 6 & \boxed{10} & 8
\end{array}
$$

In the first row we box the left-to-right maxima to the right of 1 that are not right-to-left minima. Here, those are 7 and 9. In the second row we insert a new largest letter, 10, immediately to the left of 1 and box it. Finally, in the third row, we cyclically shift the sequence of boxed letter one step to the left, thus obtaining π'. Let us call this map β.

The induced map Φ, between 231- and 321-avoiding permutations is then formally defined by

$$\Phi(\epsilon) = \epsilon; \quad \Phi(\alpha(\sigma)) = \beta(\Phi(\sigma)); \quad \Phi(\sigma \oplus \tau) = \Phi(\tau) \oplus \Phi(\sigma).$$

As an example, consider the permutation 5213476 in $\mathcal{S}_6(231)$. Decompose it using $\oplus$ and α:

$$5213476 = 52134 \oplus 21 = \alpha(2134) \oplus \alpha(1) = \alpha(\alpha(1) \oplus 1 \oplus 1) \oplus \alpha(1).$$

Reverse the order of summands and change each α to β:

$$\beta(1) \oplus \beta(1 \oplus 1 \oplus \beta(1)) = 21 \oplus \beta(1243) = 21 \oplus 41253 = 2163475.$$

In conclusion, $\Phi(5213476) = 2163475$.

5. Proof of Theorem 2

In the following five subsections we prove Theorem 2—one subsection for each equality in the theorem.

5.1. **Simion-Schmidt versus Φ.** We prepare for this proof by characterizing the Simion-Schmidt bijection in terms of left-to-right minima. (That characterization can be said to be implicit in [14].) We also characterize the bijection Φ in terms of right-to-left minima.

Definition 4. *For a permutation $\pi = a_1 a_2 \dots a_n$ of length n, define*

$$\mathfrak{lmin}(\pi) = \big\{ (i, a_i) \mid a_i \text{ is a left-to-right minima in } \pi \big\}$$

as the set of positions of left-to-right minima together with their values. Also, define $\mathcal{S}_n/\mathfrak{lmin}$ as the set of equivalence classes with respect to the equivalence induced by $\mathfrak{lmin}$: that is, π is equivalent to τ if $\mathfrak{lmin}(\pi) = \mathfrak{lmin}(\tau)$. Similarly, define $\mathfrak{rmin}$, $\mathfrak{lmax}$ and $\mathfrak{rmax}$.

The cardinality of $\mathfrak{lmin}(\mathcal{S}_n) = \{\mathfrak{lmin}(\pi) \mid \pi \in \mathcal{S}_n\}$ is easily seen to be C_n, the n-th Catalan number. The following lemma strengthens that observation:

Lemma 5. *Each equivalence class in $\mathcal{S}_n/\mathfrak{lmin}$ contains exactly one permutation that avoids 123 and one that avoids 132. In other words, both $\mathcal{S}_n(123)$ and $\mathcal{S}_n(132)$ are complete sets of representatives for $\mathcal{S}_n/\mathfrak{lmin}$.*

Proof. Given a set L in $\mathfrak{lmin}(\mathcal{S}_n)$, this is how we construct the corresponding permutation $\tau = c_1 c_2 \dots c_n$ in $\mathcal{S}_n(132)$: For i from 1 to n, if (i, a) is in L let $c_i = a$; otherwise, let c_j be the smallest letter not used that is greater than all the letters used thus far.

Given a set L in $\mathfrak{lmin}(\mathcal{S}_n)$, this is how we construct the corresponding permutation $\pi = a_1 a_2 \dots a_n$ in $\mathcal{S}_n(123)$: For i from 1 to n, if (i, c) is in L let $a_i = c$; otherwise, let a_j be the largest letter not used thus far.

It is easy to see that filling in the letters in any other way than the two ways described will either change the sequence of left-to-right minima or result in an occurrence of 132 or 123. $\qquad\square$

As an illustration of the preceding proof, with $L = \{(1,6), (3,3), (4,2), (6,1)\}$ we get 67324158 in $\mathcal{S}_8(132)$ and 68327154 in $\mathcal{S}_8(123)$.

Using Lemma 5 we can thus define a bijection between $\mathcal{S}_n(123)$ and $\mathcal{S}_n(132)$ by letting π correspond to σ if $\mathfrak{lmin}(\pi) = \mathfrak{lmin}(\sigma)$. However, this map is not new—it is the Simion-Schmidt bijection:

Lemma 6. *For π in $\mathcal{S}_n(123)$ and σ in $\mathcal{S}_n(132)$, the following two statements are equivalent:*

    (1) Simion-Schmidt$(\pi) = \sigma$;
    (2) $\mathfrak{lmin}(\pi) = \mathfrak{lmin}(\sigma)$.

Indeed, looking at the algorithm defining the Simion-Schmidt bijection we see that the variable x keeps track of the smallest letter read thus far; lines 3 and 4 express that left-to-right minima are left unchanged; and line 6 assign c_j to be the smallest letter not used that is greater than all the letters used thus far (as described above).

Here is a characterization of Φ in terms of $\mathfrak{rmin}$:

Lemma 7. *For π in $\mathcal{S}_n(231)$ and σ in $\mathcal{S}_n(321)$, the following two statements are equivalent:*

    (1) $\Phi(\pi) = \sigma$;
    (2) $(n + 1 - i, a) \in \mathfrak{rmin}(\pi) \iff (n + 1 - a, i) \in \mathfrak{rmin}(\sigma)$.

Proof. That the latter statement characterizes a bijection from $\mathcal{S}_n(231)$ to $\mathcal{S}_n(321)$ follows from Lemma 5, so all we need to show is that Φ is that bijection.

We use induction on n, the length of the permutation. The case $n = 1$ is obvious: the right-to-left minimum $(1, 1)$ goes to the right-to-left minimum $(1, 1)$. For the

induction step we distinguish two cases: π is decomposable and π is indecomposable (see Section 4 for definitions).

Suppose that π is indecomposable, and hence $\pi = \alpha(\sigma)$ for some σ. By definition, $\Phi(\pi) = \beta(\Phi(\sigma))$. The claim follows from the following equivalences:

$$
\begin{aligned}
(n + 1 - i, a) \in \mathfrak{rmin}(\pi) &\iff (n - i, a) \in \mathfrak{rmin}(\sigma) \\
&\iff (n - a, i) \in \mathfrak{rmin}\,\Phi(\sigma) \\
&\iff (n + 1 - a, i) \in \mathfrak{rmin}\,\Phi(\pi).
\end{aligned}
$$

Here, the first equivalence is immediate from the definition of α—recall that all α does is to insert a new largest letter in front of σ. The second equivalence holds by induction. The third equivalence follows from the definition of β: the cyclic shift involves no right-to-left minima and the space for the new letter is created to the left of the letter 1; therefore, 1 is added to the indices of right-to-left minima.

Suppose $\pi = \tau \oplus \rho$ is decomposable, and let $k = |\tau|$ and $\ell = |\rho|$. By definition, $\Phi(\pi) = \Phi(\rho) \oplus \Phi(\tau)$. We have

$$
\begin{aligned}
(n + 1 - i, a) &\in \mathfrak{rmin}(\pi) \\
&\iff (n + 1 - i, a) \in \mathfrak{rmin}(\tau) \text{ or } (\ell + 1 - i, a - k) \in \mathfrak{rmin}(\rho) && \text{def' of } \oplus \\
&\iff (k + 1 - (i - \ell), a) \in \mathfrak{rmin}(\tau) \text{ or } (\ell + 1 - i, a - k) \in \mathfrak{rmin}(\rho) && n = k + \ell \\
&\iff (k + 1 - a, i - \ell) \in \mathfrak{rmin}\,\Phi(\tau) \text{ or } (\ell + 1 - (a - k), i) \in \mathfrak{rmin}\,\Phi(\rho) && \text{induction} \\
&\iff (k + 1 - a, i - \ell) \in \mathfrak{rmin}\,\Phi(\tau) \text{ or } (n + 1 - a, i) \in \mathfrak{rmin}\,\Phi(\rho) && n = k + \ell \\
&\iff (n + 1 - a, i) \in \mathfrak{rmin}(\Phi(\rho) \oplus \Phi(\tau)) && \text{def' of } \oplus
\end{aligned}
$$

from which the claim follows. □

We now turn to the proof of the first identity of Theorem 2. It is equivalent to

$$\text{inverse} \circ \text{Simion-Schmidt} \circ \text{reverse} \circ \Phi \circ \text{reverse} = \text{identity}.$$

With all the preparation we have done, this is easy to prove:

$$
\begin{aligned}
(i, a) \in \mathfrak{lmin}(\pi) &\iff (n + 1 - i, a) \in \mathfrak{rmin}\,.\,\text{reverse}(\pi) \\
&\iff (n + 1 - a, i) \in \mathfrak{rmin}\,.\,\Phi.\,\text{reverse}(\pi) \\
&\iff (a, i) \in \mathfrak{lmin}\,.\,\text{reverse}.\Phi.\,\text{reverse}(\pi) \\
&\iff (a, i) \in \mathfrak{lmin}\,.\,\text{Simion-Schmidt}\,.\,\text{reverse}.\Phi.\,\text{reverse}(\pi) \\
&\iff (i, a) \in \mathfrak{lmin}\,.\,\text{inverse}\,.\,\text{Simion-Schmidt}\,.\,\text{reverse}.\Phi.\,\text{reverse}(\pi).
\end{aligned}
$$

5.2. Simion-Schmidt versus Krattenthaler.
In Lemma 6 we characterized the Simion-Schmidt bijection. We shall do the same for Krattenthaler's bijection. We start by looking at the standard bijection from 132-avoiding permutations to Dyck paths (as defined in Section 3.1).

Let P be a Dyck path of length $2n$; index its up- and down-steps 1 through n. For instance,

$$P = u_1 u_2 u_3 d_1 d_2 u_4 u_5 d_3 d_4 u_6 d_5 d_6.$$

A *peak* in a Dyck path is an up-step directly followed by a down-step. Define

$$\mathfrak{peak}(P) = \{\, (i, j) \mid u_i d_j \text{ is a peak in } P \,\}.$$

For instance, with P as before, we have $\mathfrak{peak}(P) = \{(3, 1), (5, 3), (6, 5)\}$.

Lemma 8. *Let f be the standard bijection from $\mathcal{S}_n(132)$ to $\mathcal{D}_n$. For π in $\mathcal{S}_n(132)$ and P in $\mathcal{D}_n$, the following two statements are equivalent:*

(1) $f(\pi) = P$;
(2) $(i, n + 1 - a) \in \mathfrak{lmin}(\pi) \iff (a, i) \in \mathfrak{peak}(P)$.

Proof. Clearly, knowing $\mathfrak{peak}(P)$ is equivalent to knowing the path P. Thus the second statement determines a bijection (by Lemma 5). It remains to show that the first statement implies the second.

We shall use induction on the length of the permutation. Assume that π^r is indecomposable (with respect to $\oplus$). It is easy to see that π ends with its largest letter. Hence, $\pi = \tau n$ for some τ in $\mathcal{S}_{n-1}(132)$. Let $Q = f(\tau)$. Then $P = f(\pi) = uf(\tau)d = uQd$ and

$$
\begin{aligned}
(i, n + 1 - a) \in \mathfrak{lmin}(\pi) &\iff (i, n + 1 - a) = (i, n - (a - 1)) \in \mathfrak{lmin}(\tau) \\
&\iff (a - 1, i) \in \mathfrak{peak}(Q) \\
&\iff (a, i) \in \mathfrak{peak}(P).
\end{aligned}
$$

Assume that π^r is decomposable, so $\pi^r = \rho^r \oplus \tau^r$ for some τ in $\mathcal{S}_k(132)$ and ρ in $\mathcal{S}_\ell(132)$ with both k and ℓ positive and $k + \ell = n$. Let $Q = f(\tau)$ and $R = f(\rho)$. Then $P = f(\pi) = f(\tau)f(\rho) = QR$ and

$$
\begin{aligned}
(i, &n + 1 - a) \in \mathfrak{lmin}(\pi) \\
&\iff (i, k + 1 - a) \in \mathfrak{lmin}(\tau) \text{ or } (i - k, n + 1 - a) \in \mathfrak{lmin}(\rho) \\
&\iff (i, k + 1 - a) \in \mathfrak{lmin}(\tau) \text{ or } (i - k, \ell + 1 - (a - k)) \in \mathfrak{lmin}(\rho) \\
&\iff (a, i) \in \mathfrak{peak}(Q) \text{ or } (a - k, i - k) \in \mathfrak{peak}(R) \\
&\iff (a, i) \in \mathfrak{peak}(P),
\end{aligned}
$$

which completes the proof. $\qquad\qquad\square$

Lemma 9. *Let K be Krattenthaler's bijection from $\mathcal{S}_n(123)$ to $\mathcal{D}_n$ as described in Section 3.6. For π in $\mathcal{S}_n(123)$ and P in $\mathcal{D}_n$, the following two statements are equivalent:*

(1) $K(\pi) = P$;
(2) $(i, n + 1 - a) \in \mathfrak{rmax}(\pi) \iff (a, i) \in \mathfrak{peak}(P)$.

Proof. This is an easy consequence of K's definition. $\qquad\qquad\square$

Putting Lemma 8 and Lemma 9 together we get the desired characterization of Krattenthaler's bijection.

Lemma 10. *For π in $\mathcal{S}_n(123)$ and σ in $\mathcal{S}_n(132)$, the following two statements are equivalent:*

(1) Krattenthaler$(\pi) = \sigma$;
(2) $(n + 1 - i, a) \in \mathfrak{rmax}(\pi) \iff (n + 1 - a, i) \in \mathfrak{lmin}(\sigma)$.

Having established this characterization, the rest is easy. From the sequence of equivalences

$(i, a) \in \mathfrak{lmin}(\pi)$

$\iff (n + 1 - i, a) \in \mathfrak{rmin} . \mathrm{reverse}(\pi)$

$\iff (a, n + 1 - i) \in \mathfrak{lmax} . \mathrm{inverse} . \mathrm{reverse}(\pi)$

$\iff (n + 1 - a, n + 1 - i) \in \mathfrak{rmax} . \mathrm{reverse} . \mathrm{inverse} . \mathrm{reverse}(\pi)$

$\iff (i, a) \in \mathfrak{lmin} . \mathrm{Krattenthaler} . \mathrm{reverse} . \mathrm{inverse} . \mathrm{reverse}(\pi)$

$\iff (i, a) \in \mathfrak{lmin} . \mathrm{Simion\text{-}Schmidt}^{-1} . \mathrm{Krattenthaler} . \mathrm{reverse} . \mathrm{inverse} . \mathrm{reverse}(\pi)$

it follows that

$$\mathrm{Simion\text{-}Schmidt}^{-1} \circ \mathrm{Krattenthaler} \circ \mathrm{reverse} \circ \mathrm{inverse} \circ \mathrm{reverse} = \mathrm{identity}$$

as desired.

5.3. **Knuth-Richards versus** Φ. Consider the Dyck path $P = uudduududuuddd$ of semi-length $n = 7$. Let us index its up- and down-steps 1 through n:

$$P = u_1 u_2 d_1 d_2 u_3 u_4 d_3 u_5 d_4 u_6 u_7 d_5 d_6 d_7.$$

From this path we shall construct a permutation $\pi = a_1 a_2 \ldots a_n$. Scan P's down-steps from left to right: if d_i is preceded by an up-step u_j, then let $a_{n+1-j} = i$; otherwise, let j be the largest value for which a_j is unset, and let $a_j = i$. Like this:

(1) d_1 is preceded by the up-step u_2; let $a_{8-2} = a_6 = 1$.
(2) d_2 is preceded by a down-step; let $j = 7$ and $a_7 = 2$.
(3) d_3 is preceded by the up-step u_4; let $a_{8-4} = a_4 = 3$.
(4) d_4 is preceded by the up-step u_5; let $a_{8-5} = a_3 = 4$.
(5) d_5 is preceded by the up-step u_7; let $a_{8-7} = a_1 = 5$.
(6) d_6 is preceded by a down-step; let $j = 5$ and $a_5 = 6$.
(7) d_7 is preceded by a down-step; let $j = 2$ and $a_2 = 7$.

The resulting permutation is $\pi = 5743612$. What we have just described is the algorithm defining Richard's bijection. Lines 3, 4 and 5 of that algorithm covers the case when d_i is preceded by an up-step; lines 6, 7 and 8 the case when d_i is preceded by a down-step.

Plainly, if d_i is preceded by an up-step u_i then $u_i d_j$ is a peak in P. Moreover, $a_{n+1-j} = i$ is a left-to-right minimum in the corresponding permutation. To be precise we have the following lemma.

Lemma 11. *Let* Richards *be Richards' bijection from* $\mathcal{S}_n(123)$ *to* $\mathcal{D}_n$ *as described in Section 3.4. For π in $\mathcal{S}_n(123)$ and P in $\mathcal{D}_n$, the following two statements are equivalent:*

(1) $\mathrm{Richards}(P) = \pi$;
(2) $(n + 1 - i, a) \in \mathfrak{lmin}(\pi) \iff (i, a) \in \mathfrak{peak}(P)$.

Using Lemma 8 we get a characterization of the Knuth-Richards bijection:

Lemma 12. *For π in $\mathcal{S}_n(132)$ and σ in $\mathcal{S}_n(123)$, the following two statements are equivalent:*

(1) $\mathrm{Knuth\text{-}Richards}(\pi) = \sigma$;
(2) $(i, a) \in \mathfrak{lmin}(\pi) \iff (a, i) \in \mathfrak{lmin}(\sigma)$.

The rest is easy. We have

$$(i, a) \in \mathfrak{lmin}(\pi) \iff (a, i) \in \mathfrak{lmin} . \text{Knuth-Richards}^{-1}(\pi)$$
$$\iff (n + 1 - a, i) \in \mathfrak{rmin} . \text{reverse} . \text{Knuth-Richards}^{-1}(\pi)$$
$$\iff (n + 1 - i, a) \in \mathfrak{rmin} . \Phi . \text{reverse} . \text{Knuth-Richards}^{-1}(\pi)$$
$$\iff (i, a) \in \mathfrak{lmin} . \text{reverse} . \Phi . \text{reverse} . \text{Knuth-Richards}^{-1}(\pi)$$

and hence

$$\text{reverse} \circ \Phi \circ \text{reverse} \circ \text{Knuth-Richards}^{-1} = \text{identity} .$$

5.4. **Simion-Schmidt versus Mansour-Deng-Du.** We will show that

$$\text{Mansour-Deng-Du} = \text{reverse} \circ \text{Simion-Schmidt} \circ \text{reverse} .$$

Due to Lemma 6 it suffices to prove this lemma:

Lemma 13. *For π in $\mathcal{S}_n(132)$ and σ in $\mathcal{S}_n(123)$, the following two statements are equivalent:*

(1) *Mansour-Deng-Du$(\pi) = \pi'$;*
(2) $\mathfrak{rmin}(\pi) = \mathfrak{rmin}(\pi')$.

Proof. Assume that π and π' are as above. According to the proofs of Corollaries [9, Cor. 4] and [9, Cor. 16], the positions of right-to-left minima in π and π' are the same and, in particular, $\text{rmin}(\pi) = \text{rmin}(\pi')$. Thus we only need to prove that right-to-left minima are preserved in value under the Mansour-Deng-Du bijection. Equivalently, we need to prove that non-right-to-left-minima (n-rmin) are preserved in value.

One can see that a letter a is an n-rmin in π if and only if the reduced decomposition of π contains a run of simple transpositions $(s_{a-1} \ldots)$. In particular, $a = 1$ is always an n-rmin. Thus $\pi = (12 \ldots n)\sigma_1 \ldots \sigma_k$ and $\pi' = (12 \ldots n)\sigma'_1 \ldots \sigma'_k$ for $k = n - \text{rmin}(\pi)$. That is, π and π' have the same number of runs of simple transpositions in the reduced decompositions and it remains to show that the first letter of σ_j equals the first letter of σ'_j whenever $1 \leq j \leq k$.

Let P be the intermediate Dyck path and consider its $(x + y)$- and $(x - y)$-labellings of P. Note that cells touching the x-axis receive the same labels under both labellings. From this, and the way that the zigzag and trapezoidal decompositions are constructed, it immediately follows that σ_k and σ'_k begin with the same letter, namely, the label C of the rightmost cell.

We now proceed by induction on the number of essential cells. If there are no essential cells, then the statement is true. Suppose we have $k > 0$ essential cells. Remove the rightmost zigzag strip to get a Dyck path P'. Note that $|P'| = |P| - 2$ and that P' has $k - 1$ essential cells. Clearly, the permutation corresponding to the $(x + y)$-labeling of P' is $\tau = (12 \ldots (n - 1))\sigma_1 \ldots \sigma_{k-1}$. Let the permutation corresponding to the $(x - y)$-labeling of P' be $\tau'' = (12 \ldots (n-1))\sigma''_1 \ldots \sigma''_{k-1}$. From the properties of the $(x - y)$-labeling, a cell labeled $Q \neq C$ is the rightmost cell of a trapezoidal strip in P if and only if Q is the rightmost cell of a trapezoidal strip in P'. This means that σ'_i and σ''_i begin with the same letter for $1 \leq i \leq k - 1$. The desired result now follows from the induction hypothesis applied to P', τ and τ''. $\qquad\square$

FIGURE 3.

5.5. Reifegerste versus Knuth-Rotem. We will show that

$$(1) \qquad \text{Reifegerste} = \text{inverse} \circ \text{Knuth-Rotem} \circ \text{inverse}.$$

Reifegerste's bijection is defined by values and positions of excedances. Suppose that $\pi = a_1 a_2 \ldots a_n$ is a 321-avoiding permutation and that π' is the image of π under Reifegerste's bijection. The only permutation of length n without excedances is $12 \ldots n$. As is easy to see, that permutation is fixed by both sides of identity (1). So we can assume that π has at least one excedance. Say that has an excedance a at position i, denoted (i, a). Also, let (j, b) be the excedance closest to the left of (i, a). If no such excedance exists we define $j = 0$. Note that $b < a$, since otherwise an occurrence of 321 would be formed.

Consider the Ferrer's diagram corresponding to π in the definition of Reifegerste's bijection (shaded in the figure in Subsection 3.7 on page 10). The point $(j+1, n+2-a)$ is a corner in that diagram. It is sent to the left-to-right minimum $(j+1, n+2-a)$ in π'. Thus, we need to prove that an excedance (i, a) corresponds to a left-to-right minimum $(j + 1, n + 2 - a)$ under the right hand side of identity (1).

Because $a > i$, $b > j$ and $b < a$ it follows that both i and j are *strict non-excedances* in π^{-1} with the property that there are no other strict non-excedances between i and j. (A strict non-excedance between i and j would either bring an occurrence of 321 in π or in π^{-1}, or an occurrence of an excedance between a and b in π.) Thus the ballot sequence β, obtained when applying Knuth-Rotem's bijection to π^{-1} will have the letter $j + 1$ in positions $b = a_j, a_{j+1}, \ldots, a_{i-1}$ and the letter $i + 1$ in position $a = a_i$. Let P be the Dyck path corresponding to β. It remains to show that f^{-1}, the inverse of the standard bijection, sends P to a permutation having the letter $j + 1$ in position $n + 2 - a$. After applying inverse we would then have the letter $n + 2 - a$ in position $j + 1$, the same outcome as when applying Reifegerste's bijection.

For the remainder of this proof we use induction on n, the length of the permutation π. The smallest permutation that have an excedance is $\pi = 21$. In this case, $i = 1$, $j = 0$ and $a = 2$. See Figure 3A. After rotating that diagram counter clockwise by $3\pi/4$ radians we read the Dyck path $udud$. The inverse of the standard bijection, f^{-1}, sends $udud$ to the permutation $\pi' = 21$. It has the letter $j + 1 = 1$ in position $n + 2 - a = 2$ as desired.

Assume that $n > 2$. Let D be the diagram constructed from the ballot-sequence corresponding to π. Let P be the Dyck path we read from D. Let (r, s) be the coordinate in D corresponding to the first return to the x-axis in P. In particular, $r = s$. Consider the following four cases.

Case 1, $s = 0$. This case is sketched in Figure 3B. Note that $a > i + 1$, and we can remove the first step and the last step of the path, thus shifting the path one step to the left. This corresponds to adjoining n to the right side of the permutation obtained inductively from the reduced Dyck path. Moreover, the number of steps is now $2(n-1)$, i and j are unchanged, while a becomes $a - 1$. From the induction hypothesis it follows that $j + 1$ will be in position $(n-1) + 2 - (a-1) = n + 2 - a$, as desired.

Case 2, $n > s > i$. From the definition of the standard bijection we see that the permutation $f^{-1}(P)$ is of the form $\pi' = \sigma n \tau$ where each letter of σ is larger than any letter of τ; the largest letter, n, is in position $n - s + 1$; part σ is obtained inductively from the portion of the path above the line $y = s$; and τ is obtained inductively from the portion of the path below the line $y = s$. Applying the induction to τ we see that the letter $j + 1$ is in position $(s - 1) + 2 - a = s + 1 - a$. Thus, in π', the letter $j + 1$ is in position $(s + 1 - a) + (n - s + 1) = n + 2 - a$.

Case 3, $0 < s < i$. From the path we remove the part below the line $y = s$, we remove the first and the last steps of the remaining path, and we shift the obtained path so that it goes from $(n - s - 1, n - s - 1)$ to $(0,0)$. The resulting path we call P'. It is responsible for building the standardization σ' of σ in the outcome permutation $\pi' = \sigma n \tau$. (Here σ' is obtained from σ by adding $|\tau|$ to each of its letters.) Note that i, j and a in P become $i - s$, $j - s$ and $a - s - 1$ in P', respectively. From the induction hypothesis applied to P' it follows that $(j - s) + 1$ will be in position $(n - s - 1) + 2 - (a - s - 1) = n + 2 - a$ in σ'. But $(j - s) + 1$ in σ' is $(j - s) + 1 + s = j + 1$ in σ. Thus, in π', the letter $j + 1$ will be in position $n + 2 - a$.

Case 4, $s = i$. In this case $i + 1 = a$ and we need to prove that the letter $j + 1$ is in position $n + 2 - a = n + 1 - i$. Remove from P the first and last steps of the subpath from (n, n) to (i, i). Shift the path from $(n, n - 1)$ to $(i + 1, i)$ one step to the left. Let P' denote the obtained path from $(n - 1, n - 1)$ to $(0,0)$. Now we apply f^{-1} to P and get a 132-avoiding permutation of length $(n - 1)$ of the form $\sigma (n - 1) \tau$, where $|\tau| \geq i$. The first return in P' to the line $y = x$ will be either at (i, i) or above the line $y = i$. By the induction hypothesis, in $\sigma(n-1)\tau$, the letter $j + 1$ is in position $(n - 1) + 2 - a = n - i$. This position is located to the right of $n - 1$. Looking at the definition of the standard bijection we see that $f^{-1}(P)$ is $\sigma(n-1)\tau$ with the largest letter, n, inserted immediately to the left of the letter $j + 1$, causing it to appear in position $n - i + 1 = n + 2 - a$.

6. Proof of Theorem 1

In this section we prove the equidistributions presented in Theorem 1. In light of Theorem 2 it suffices to consider the following five bijections:

$$\text{Simion-Schmidt, Reifegerste, West, Knuth and Elizalde-Deutsch}.$$

For brevity we shall in this section write $\text{stat}_1 \simeq \text{stat}_2$ when, for all π in the relevant domain, we have $\text{stat}_1(\pi) = \text{stat}_2(\psi(\pi))$, where ψ is the bijection under consideration.

6.1. **Simion-Schmidt's bijection.** Let π be a 123-avoiding permutation of length n and let π' be the image of π under the Simion-Schmidt bijection. By Lemma 6 we have $\mathrm{lmin}(\pi) = \mathrm{lmin}(\pi')$. In particular, lmin $\simeq$ lmin. A three letter segment abc is a valley precisely when a and b are left-to-right minima but c is not. Thus valley $\simeq$ valley. Similarly, we have ldr $\simeq$ ldr and head $\simeq$ head. Indeed, ldr is

determined by the position of the first non-left-to-right minimum and head is the first left-to-right minimum.

- comp.r $\simeq$ comp.r: Suppose that $\pi = AB$ where B is the rightmost reverse component. Then $B = mC$ where m is a left-to-right minimum in π. It follows that $\pi' = A'B'$ in which A' and B' are the images, under Simion-Schmidt's bijection, of A and B, respectively. Because the Simion-Schmidt bijection is an involution it also follows that B' is the rightmost reverse component in π'. By induction on the number of reverse components we thus have comp.r $\simeq$ comp.r.

- slmax.c $\simeq$ lir: The statistic slmax.c is the position of the second left-to-right minimum (it is defined to be n if there is only one left-to-right minimum). Suppose $\pi = a_1 A a_2 B$ where a_1 and a_2 are the two leftmost left-to-right minima (the case $\pi = 1 \ldots n$ is trivial). Note that each letter in A is larger than a_1 and a_2. We know that $\pi' = a_1 A' a_2 B'$ and $|A| = |A'|$. To avoid the pattern 132 the segment A' must be increasing. Thus slmax.c $\simeq$ lir.

- slmax.i.r $\simeq$ lir.i: The statistic slmax.i.r is one less than the minimal i such that the letter i is to the left of the letter 1. Note that such an i in π must be a left-to-right minimum. Suppose slmax.i.r$(\pi) = i-1$ and $i \leq n$ (the case slmax.i.r$(\pi) = n$ is trivial). Then i and 1 are two consecutive left-to-right minima in π. Consequently, i and 1 are two consecutive left-to-right minima in π'. Thus the letters $2, 3, \ldots, i-1$ must be to the right of 1 in π'. To avoid forming an occurrence of 132, those letters must also be in increasing order. Thus lir.i$(\pi') = i - 1$.

- ldr.i $\simeq$ ldr.i: By definition ldr.i(π) is the largest i such that $i, i - 1, \ldots, 1$ is a subword in π. In particular, $i + 1$, if it exists, is to the right of i in π. Suppose that ldr.i$(\pi) = i$. Clearly, the letters $i, i - 1, \ldots, 1$ are consecutive left-to-right minima in π and $i + 1$ is a non-left-to-right minimum. Thus $i, i - 1, \ldots, 1$ is a subword in π' and $i + 1$ is to the right of i in π'; hence ldr.i$(\pi') = i$.

- head.i.r $\simeq$ rmin: Plainly, head.i.r$(\pi) = n - i + 1$ where i is the position of 1 in π. Let $\pi' = \sigma 1 \tau$. The letter 1 is in the same position in π' as it is in π and so $|\tau| = n - i$. To avoid 132 the letters of τ must be in increasing order. The sequence of right-to-left minima in π' is thus simply 1τ and therefore head.i.r$(\pi) = $ rmin(π').

- valley.i $\simeq$ valley.i: By definition, valley.i(π) is the number of letters i in π such that i is to the left of both $i - 1$ and $i + 1$. Suppose i in π is one of those letters counted by valley.i(π). To avoid 123 the letter i must be a left-to-right minimum. Thus $i - 1$ is a left-to-right minimum too, but $i + 1$ is not. Since $\mathsf{lmin}(\pi) = \mathsf{lmin}(\pi')$ this observation translates to π'. That is, in π', the letters i and $i - 1$ are left-to-right minima whereas $i + 1$ is not. Thus $i + 1$ is to the right of i in π' and the letter i in π' is counted by valley.i(π'). Because the Simion-Schmidt bijection is an involution it is easy to see that no i which is not counted by valley.i(π) contributes to valley.i(π').

- rank $\simeq$ rank: Suppose that rank$(\pi) = k$ and π has a in position $k + 1$. We distinguish two cases based on whether a is a left-to-right minimum. If a is a non-left-to-right-minimum, then $k+1$ is the left-to-right minimum closest to a in π from the left. Since $\mathsf{lmin}(\pi) = \mathsf{lmin}(\pi')$ we have rank$(\pi') = k$. On the other hand, if a is a left-to-right minimum, then $a \leq k + 1$ and π' will have the same left-to-right minimum, a, in position $k + 1$. Thus rank$(\pi') = k$ in this case as well.

6.2. **Reifegerste's bijection.** Let $\pi = a_1 a_2 \ldots a_n$ be a 321-avoiding permutation of length n and let π' be the image of π under Reifegerste's bijection. That exc $\simeq$ des was proved in [10].

- valley $\simeq$ valley: A 321-avoiding permutation is a shuffle of two increasing subwords. From this one can see that if $a_i a_{i+1} a_{i+2}$ is a valley in π, then a_i is an excedance and a_{i+1} is a non-excedance. Thus, in the permutation matrix corresponding to π, there is an E-square in the i-th row but no E-square in the $(i+1)$-th row. It follows that the Ferrer's diagram has a corner in the $(i+1)$-th row. Thus $a'_i a'_{i+1} a'_{i+2}$ is a valley in $\pi' = a'_1 a'_2 \ldots a'_n$. (The dot in row $i+1$, corresponding to the letter a'_{i+1}, will be to the left of the dots in rows i and $(i+2)$, corresponding to the letters a'_i and a'_{i+2}, respectively).

- peak.i $\simeq$ valley.i: By definition, peak.i(π) is the number of letters a in π to the right of both $a-1$ and $a+1$; and valley.i(π) is the number of letters a in π to the left of both $a-1$ and $a+1$. If i is counted by peak.i, then i is a non-excedance while $i+1$ is an excedance; otherwise an occurrence of 321 is formed. Thus we have an E-square in column $n-i$, but no E-square in column $n-i+1$. Not also that there is a column $n-i+2$ (corresponding to the letter $i-1$). Thus column $n-i+1$ is not the rightmost column of the matrix and it contains a corner of the Ferrer's diagram. So peak.i counts the columns that contain a corner but no E-squares, excluding the rightmost column. In the construction of π' each corner with the properties described above gives rise to a letter i in π' to the left of $i-1$ and $i+1$. Indeed, such a corner has an E-square in the column immediately to its left and no E-square in the column immediately to its right; consequently, the points corresponding to $i-1$ and $i+1$ in π' will be below the point corresponding to i in π'. Thus peak.i $\simeq$ valley.i.

- slmax.i $\simeq$ zeil: The letter between the two leftmost left-to-right maxima in a 321-avoiding permutation is an initial segment of $123\ldots$. Hence slmax.i is the length of the maximal initial segment of the form $23\ldots i$. Moreover, each of the letters counted by slmax.i is an excedance. Thus we have E-squares in positions $(\ell, n-\ell)$ for $\ell = 1, 2, \ldots, i-1$ and no E-square in position $(k, n-k)$. Thus π' is of the form $\pi' = n(n-1)(n-2)\ldots(n-k+2)A(n-k)B(n-k+1)C$, where A, B, and C are some words. Clearly, zeil$(\pi') = k =$ slmax.i(π).

- head.i $\simeq$ ldr: The statistic head.i(π) is the position of 1 in π. Say this position is i. Then row i of the permutation matrix corresponding to π is the topmost row that does not contain an E-square. Thus, in π', the first i letters will be in decreasing order, while in position i there will be an ascent (unless $i = n$). Thus ldr$(\pi') = i$.

- slmax.r.c $\simeq$ rdr: Let the right-to-left minima of π, read from right to left, be r_1, r_2, $\ldots$, r_k. One can see that slmax.r.c is one more than the number of letters between r_1 and r_2. (To make sure there is at least two left-to-right minima we can assume that a 0 stays in front of π when considering this statistic.) To avoid the pattern 321 the letters between r_1 and r_2 must be the largest letters in π, and thus all of them are excedances. Thus, in the permutation matrix corresponding to π, there will be E-squares in positions $(i, n-i)$ for $i = n-1, n-2, \ldots, n - (\text{slmax.r.c}(\pi) - 1)$ and there will be no E-squares in the $(n - \text{slmax.r.c}(\pi))$-th row. In turn, this guarantees that the 132-avoiding permutation π' ends with $j(\text{slmax.r.c}(\pi) - 1)(\text{slmax.r.c}(\pi) - 2)\ldots 1$, where j, if it exists, is strictly larger than slmax.r.c(π). To avoid 132 this segment must be preceded by a letter smaller than j. Thus rdr$(\pi') =$ slmax.r.c(π).

- rir $\simeq$ rmin: Let rir$(\pi) = i$. We have $i = n$ only if $\pi = \pi' = 12\ldots n$, and then, trivially, rir$(\pi) =$ rmin$(\pi') = n$. Assume $i < n$. The letters in the rightmost

increasing run are non-excedances while the letter immediately to the left of that run is an excedance. Thus the bottom-most E-square is in row $n-i$ in the permutation matrix. The dots placed in rows $n-i+1$, $n-i+2,\ldots,n$ when creating π' gives the sequence of right-to-left minima in π', and thus $\mathrm{rmin}(\pi') = i$.

• lir.i $\simeq$ lmax: Let lir.i$(\pi) = i$. We have $i = n$ only if $\pi = \pi' = 12\ldots n$, and then, trivially, lir.i$(\pi) = \mathrm{lmax}(\pi') = n$. Assume $i < n$. By definition, i is the largest positive integer such that $1, 2, \ldots, i$ is a subword of π. Since π is the shuffle of two increasing sequences, $i + 1$ is the leftmost excedance in π. The E-square corresponding to $i+1$ in the permutation matrix is placed in column $n-i$, leaving i columns to the right of it. When constructing π' each of those i columns get a dot, beginning at the position $(1, n-i+1)$ and going in the South-East direction. Those dots give the sequence of left-to-right maxima in π', and thus $\mathrm{lmax}(\pi') = i$.

• last $\simeq$ m-ldr.i: Note that ldr.i(π) is the largest i such that $i, (i-1), \ldots, 1$ is a subword in π. The case $\pi = \pi' = 12\ldots n$ is trivial. Suppose $i < n$ and last$(\pi) = i$. To avoid the pattern 321, the letters $i+1$, $i+2, \ldots, n$ must form a subword of π, and clearly each of them is an excedance. Therefore the $n-i$ leftmost columns of the permutation matrix corresponding to π contain E-squares but the $(n-i+1)$-th column does not contain an E-square. From this, and the way π' is constructed, it immediately follows that $(n-i+1), (n-i), \ldots, 1$ is the longest subword of π' of the sought type. Thus m-ldr.i$(\pi') = i$.

6.3. West's bijection.

Recall that West's bijection is induced by an isomorphism between generating trees. The two isomorphic trees generate 123- and 132-avoiding permutations, respectively. However, for the purpose of this proof we shall generate 321-avoiding permutations instead of 123-avoiding ones. This change is reflected in positions of active sites: In the 123-avoiding case the active sites are all the positions to the left of the leftmost ascent and the position in between the two letters of the leftmost ascent. In the 321-avoiding case the active sites are all the positions to the right of the rightmost ascent and the position in between the two letters of the rightmost descent.

The active sites in a 132-avoiding permutation are the leftmost position and every position immediately to the right of right-to-left maxima.

All the proofs below are by induction on the length of the permutation, with easily verifiable base cases. Let π be a 321-avoiding permutation of length n. Let π' be the image of π under the modified version of West's bijection as described above.

• peak.i $\simeq$ valley.i: By definition, peak.i(π) is the number of letters a in π to the right of both $a-1$ and $a+1$; and valley.i(π') is the number of letters a in π' to the left of both $a-1$ and $a+1$.

Case 1. Assume that π ends with n. Then π' begins with n. Inserting $n+1$ to the right of n in π does not change peak.i. Inserting $n+1$ to the left of n in π' does not change valley.i. Inserting $n+1$ in any other position in π and π' increases peak.i and valley.i by 1.

Case 2. Assume that π does not end with n. To avoid the pattern 321, the letter n must be the the last letter of the rightmost descent. Thus n is not the leftmost letter in π', and to avoid the pattern 132, the letter $n-1$ must be to the left of n. It follows that inserting $n+1$ in an active site of π and π', respectively, does not change peak.i and valley.i, respectively.

• exc $\simeq$ asc: Inserting $n + 1$ at the end of π does not change the number of excedances. Similarly, inserting $n+1$ at the beginning of π', the number of ascents is not changed. In all other cases, the number of excedances in π and the number of ascents in π' is increased by 1.

• slmax.i $\simeq$ lir.i: We see that slmax.i(π) is one more than the length of the maximum initial segment of the form $234\ldots$ in π. Also, lir.i(π') is the largest i such that $12\ldots i$ is a subword of π'.

Case 1. Assume that $\pi = 234\ldots n1$. Using induction it is easy to verify that $\pi' = 12\ldots n$. Inserting $n+1$ at the end of π does not change slmax.i, and inserting $n+1$ at the beginning of π' does not change lir.i. On the other hand, inserting $n+1$ between n and 1 in π increases slmax.i by 1, and inserting $n+1$ at the end of π' increases lir.i by 1.

Case 2. Assume that $\pi \neq 234\ldots n1$, and thus $\pi' \neq 12\ldots n$. Inserting $n+1$ will not change slmax.i in π and it will not change lir.i in π'.

• slmax.r.c $\simeq$ comp: By definition slmax.r.c$(\pi) = 1$ if the letter n is not in position $n-1$, and slmax.r.c(π) is one more than the length of the maximal segment of the form $i(i+1)\ldots n$ if n is in position $n-1$.

Inserting $n+1$ at the end of π' (creating two active sites) increases the number of components by 1. The corresponding action on π (creating two active sites) is inserting $n+1$ in position n; that increases slmax.r.c by 1. Inserting $n+1$ in any other active site of π' will create an indecomposable permutation (n+1 will be to the left of 1). The corresponding operation on π places $n+1$ in a position different from n, and thus slmax.r.c will be 1.

• rir $\simeq$ rmax: A proof is straightforward from the location of active sites in π and π': the (number of) active sites in π and π' essentially give the rir- and rmax-statistics, respectively.

• lir.i $\simeq$ ldr.i: By definition lir.i(π) is the largest i such that $12\ldots i$ is a subword of π, and ldr.i(π') is the largest i such that $i(i-1)\ldots 1$ is a subword of π'.

It is easy to see, by induction, that $\pi = 12\ldots n$ corresponds to $\pi' = n(n-1)\ldots 1$. Inserting $n+1$ at the end of π and at the beginning of π' will increase lir.i and ldr.i by 1, respectively. For any other insertions and any other π and π', the statistics lir.i and ldr.i will not change.

• last $\simeq$ head: Inserting $n+1$ in π changes the last letter only if the insertion is at the end. Similarly, inserting $n+1$ in π' changes the leftmost letter only if the insertion is at the beginning.

6.4. **Knuth's bijection.** Elizalde and Pak [5] proved that exc $\simeq$ exc, fix $\simeq$ fix and lis $\simeq$ n-rank. Let π be a 321-avoiding permutation of length n and let π' be the image of π under Knuth's bijection.

• lir $\simeq$ lmax: If lir$(\pi) = i$ then the first row of the *recording* tableau begins with $1, 2, \ldots, i$, and $i+1$ (if it exists) is the leftmost element in the second row. Thus the statistic lir translates to the statistic "length of the rightmost slope" (segment of down-steps) in the corresponding Dyck path. After reflection, that statistic becomes "length of the leftmost slope" (segment of up-steps). The up-steps in the leftmost slope have corresponding down-steps such that between these steps one has a proper Dyck path. In particular, the down-step corresponding to the leftmost up-step gives the first return to x-axis, giving the position of n, the rightmost left-to-right maxima, in π'. Proceeding recursively it is easy to see that

in general the up-steps in the leftmost slope read from left to right correspond to the left-to-right maxima in π' read from right to left. This gives the desired result.

• lir.i $\simeq$ rmin: The statistic lir.i is the length of the longest subword of the form $12\ldots i$. So, if lir.i$(\pi) = i$, the first row of the *insertion* tableau begins with $1, 2, \ldots, i$, and $i + 1$ (if it exists) is the leftmost element in the second row. Thus the statistic lir.i translates to the statistic "length of the leftmost slope" in the corresponding Dyck path. After reflection, that statistic becomes "length of the rightmost slope". Returns to x-axis in the Dyck path correspond to reverse components in π'. Consider the part D' of the Dyck path between the last return and next to last return to x-axis. The first up-step of D' corresponds to the rightmost down-step in the rightmost slope, and it corresponds to the rightmost letter, say a, in π'. The letter a is the largest letter in the rightmost reverse component of π', and thus it is a right-to-left minimum. Proceeding recursively we see that the next to last down-step in the rightmost slope corresponds to the second right-to-left minimum from the right, and so on.

6.5. **Elizalde-Deutsch's bijection.** Elizalde and Deutsch [4] proved that fix $\simeq$ fix.

References

[1] J. Bandlow and K. Killpatrick, An area-to-inv bijection between Dyck paths and 312-avoiding permutations, *Elect. J. Combin.* **8** (2001) #R40.

[2] A. Claesson, S. Kitaev and E. Steingrímsson, Stack sorting, trees, and pattern avoidance, preprint.

[3] I. P. Goulden and J. West, Raney paths and a combinatorial relationship between rooted nonseparable planar maps and two-stack-sortable permutations, *J. Combin. Theory Ser. A* **75** (1996), 200–242.

[4] S. Elizalde and E. Deutsch, A Simple and Unusual Bijection for Dyck Paths and its Consequences, *Annals of Comb.* **7** (2003), 281–297.

[5] S. Elizalde and I. Pak, Bijections for refined restricted permutations, *J. Comb. Theory, Ser. A* **105**(2) (2004), 207–219.

[6] D. E. Knuth, *The art of computer programming*, Vol. I *Fundamental algorithms.* Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont, 1969.

[7] D. E. Knuth, *The art of computer programming*, Vol. III, Addison-Wesley, Reading, MA, 1973.

[8] C. Krattenthaler, Permutations with restricted patterns and Dyck paths, *Adv. Appl. Math.* **27** (2001), 510–530.

[9] T. Mansour, E. Deng and R. Du, Dyck paths and restricted permutations, *Discrete Appl. Math.* **154** (2006), 1593–1605.

[10] A. Reifegerste, A generalization of Simion-Schmidt's bijection for restricted permutations, *Electron. J. Combin.* **9** (2002), no. 2, Paper #R14, 9 pp. (electronic).

[11] D. Richards, Ballot sequences and restricted permutations, *Ars Combin.* **25** (1988), 83–86.

[12] A. Robertson, D. Saracino and D. Zeilberger, Refined Restricted Permutations, *Ann. Comb.* **6** (2003), 427–444.

[13] D. Rotem, On a correspondence between binary trees and a certain type of permutations, *Information processing letters* **4**, no. 3 (1975), 58–61.

[14] R. Simion and F. W. Schmidt, Restricted Permutations, *Europ. J. Combin.* **6** (1985), 383–406.

[15] J. West, Permutation trees and the Catalan and Schröder numbers *Discrete Math.*, **146** (1995), 247–262.

[16] D. Zeilberger: A proof of Julian West's conjecture that the number of two-stack-sortable permutations of length n is 2(3n)!/((n+1)!(2n+1)!), *Discrete Math.* **102** (1992), 85–93.

The Mathematics Institute, Reykjavik University, Kringlan 1, 103 Reykjavik, Iceland