

An Equidistribution Involving Invisible Inversions

Michael Coopman[†] and Martin Rubey^{‡*}

[†]*Department of Mathematics, University of Florida, Gainesville, FL 32601*
Email: m.coopman@ufl.edu

[‡]*Fakultät für Mathematik and Geoinformation, TU Wien, Austria*
Email: martin.rubey@tuwien.ac.at

Received: November 5, 2021, **Accepted:** April 26, 2022, **Published:** May 6, 2022
 The authors: Released under the CC BY-ND license (International 4.0)

ABSTRACT: We provide two explicit bijections demonstrating that, among permutations, the number of invisible inversions is equidistributed with the number of occurrences of the vincular pattern 13-2 after sorting the set of runs.

Keywords: Bijections; Invisible inversions; Permutation statistics

2020 Mathematics Subject Classification: 05A05; 05A19

1. Introduction

The set of *visible inversions*, a certain subset of the usual inversion set of a permutation, was introduced by Z. Hamaker, E. Marberg, and P. Pawłowski in [4] to describe the rank function of the restriction of the Bruhat order to involutions.

In this note, we are concerned with the complementary set of *invisible inversions*[†] $\text{Inv}(\pi)$ of a permutation π . This is the set of pairs (i, j) such that $i < j$ and $\pi(i) > \pi(j) > i$. Querying the database of combinatorial statistics www.findstat.org suggested that the number of invisible inversions is equidistributed with the number of occurrences of the vincular pattern 13-2 after sorting the set of runs. Our goal is to prove a refinement of this guess. We note that the map sorting the set of runs of a permutation appears to have some remarkable properties, as recently demonstrated by P. Alexandersson and O. Nabawanda in [1].

An *ascending run* of π is a maximal consecutive subsequence of π in its one-line notation that is increasing. Let $\text{runsort}(\pi)$ be the permutation obtained by rearranging the set of ascending runs in increasing order of their minimal elements[‡]. Furthermore, let $13\text{-}2(\pi)$ [§] be the set of pairs (i, j) such that $i < j$ and $\pi(i) < \pi(j) < \pi(i+1)$. Then, it turns out that

$$\sum_{\pi \in \mathfrak{S}_n} q^{|\text{Inv}(\pi)|} = \sum_{\pi \in \mathfrak{S}_n} q^{|13\text{-}2(\text{runsort}(\pi))|}.$$

In fact, this result can be refined significantly. To state the refinement, let $31\text{*}2(\pi)$ be the set of pairs (i, j) such that $\pi(i) > \pi(j) > \pi(i+1)$ and the minimal element of the descending run containing $\pi(i)$ and $\pi(i+1)$ is smaller than the minimal element of the descending run containing $\pi(j)$. In this case we call $\pi(j)$ a *descent view*.

Thus, $|31\text{*}2(\pi)| = |13\text{-}2(\text{runsort}(\text{reverse}(\pi)))|$. Furthermore, recall that a *global ascent* of a permutation π is an index m such that π restricted to $[m]$ is a permutation. Finally, a *left-to-right maximum* of π is a value $v = \pi(j)$ such that $\pi(i) < v$ for all $i < j$.

Theorem 1.1. *There is an explicit bijection $\chi : \mathfrak{S}_n \rightarrow \mathfrak{S}_n$ such that, setting $\sigma = \chi(\pi)$,*

- *the multiset $\{\{\pi(j) \mid (i, j) \in 31\text{*}2(\pi)\}\}$ of descent views is the multiset $\{\{\sigma(j) \mid (i, j) \in \text{Inv}(\sigma)\}\}$ of invisible inversion bottoms,*

*Martin Rubey was supported by the Austrian Science Fund (FWF): P 29275.

[†]www.findstat.org/St001727

[‡]www.findstat.org/Mp00223

[§]www.findstat.org/St000356

- the set of maximal elements of the descending runs of π is the set of positions of the weak deficiencies of σ ,
- the set of minimal elements of the descending runs of π is the set of values of the weak deficiencies of σ ,
- the sets of global ascents of π and σ are the same, and
- the set of left-to-right maxima of π is the set of maximal elements in the cycles of σ .

We will show in Section 3 that $\chi(\pi)$ is uniquely determined on its set of exceedances by the requirements of Theorem 1.1. We will then present two slightly different ways to define $\chi(\pi)$ on the set of weak deficiencies, the first variant, www.findstat.org/Mp00235, in Section 4.1, and the second variant, www.findstat.org/Mp00237, in Section 4.2.

We remark that, for both bijections, the preimage of the set of involutions is the set of permutations whose descending runs all have length of at most two and increasing maximal elements. Restricting χ to this set we obtain the classical (inverse) fundamental transformation[¶] due to Rényi, and Foata and Schützenberger.

It may be interesting to explore whether further bijections satisfying some of the requirements of Theorem 1.1 can be found. One might speculate that encoding permutations in terms of labelled Motzkin paths with two types of horizontal steps, using the Françon-Viennot or the Foata-Zeilberger bijection, or one of its variants (see, for example, [2, sec. 5]), could be helpful. Doing so, the exceedances of $\chi(\pi)$ and their values would carry the same information as the Motzkin path itself and its labels on the up steps and the horizontal steps of one kind.

One variant of these bijections, between involutions and Motzkin paths with only one kind of horizontal steps and labels only on the down steps were in fact used by Z. Hamaker and the first author [3] to give a visual interpretation of the number of visible inversions in involutions.

2. Definitions

Let π be a permutation. A (*descending*) *run* of π is a maximal consecutive decreasing subsequence of π in its one-line notation. From now on, we will only consider descending runs. We denote the set of runs of π with $R(\pi)$. A *run top/run bottom* is the largest/smallest element of a run. In the examples we will, whenever appropriate, separate the runs of a permutation with a little space to increase readability. We recall from the introduction that a *descent view* is a value $\pi(j)$ such that there exists a descent at some position i with $\pi(i) > \pi(j) > \pi(i + 1)$ and such that the minimal element of the descending run containing $\pi(i)$ is smaller than the minimal element of the descending run containing $\pi(j)$.

A *partial permutation* is an injective map $S \rightarrow [n]$ defined on a subset S of $[n]$. An *exceedance/weak deficiency* of a (possibly partial) permutation π is a position $i \in [n]$ such that $\pi(i)$ is strictly greater/weakly less than i . An *inversion* of a (possibly partial) permutation π is a pair (i, j) such that $i < j$ and $\pi(i) > \pi(j)$. We call $\pi(j)$ the *inversion bottom*. An *invisible inversion* additionally satisfies $i < \pi(j)$.




Figure 1: Descent views of $\pi = 529687314$ and 873152496 , and invisible inversions of $\sigma_e = 375??89??$.

To display a partial permutation $\pi : S \rightarrow [n]$ we will sometimes use a one-line notation by setting $\pi(i) = ?$ for $i \in [n] \setminus S$ and writing down the values $\pi(1), \dots, \pi(n)$ of the so extended function. Graphically, we illustrate a (partial) permutation by drawing a grid of cells, with n rows and n columns, and placing, for each $i \in S$, a dot into the cell in the i -th column from the left and the $\pi(i)$ -th row from the bottom. Whenever appropriate, we separate the descending runs of π by heavier vertical lines.

Let us illustrate these definitions on some concrete examples. The permutation $\pi = 529687314$ is displayed in the left most diagram of Figure 1. Its run bottoms are $RB = \{1, 2, 4, 6\}$, its run tops are $RT = \{4, 5, 8, 9\}$ and its multiset of descent views is $\{\{2, 4, 4, 5, 6\}\}$.

[¶]www.findstat.org/Mp00086

The permutation $\bar{\pi}$ in the middle diagram is obtained by sorting the set of descending runs of π such that the sequence of run bottoms increases. It is clear that the sets of run bottoms and also the multisets $\{\{\pi(j) \mid (i, j) \in 31*2(\pi)\}\}$ and $\{\{\bar{\pi}(j) \mid (i, j) \in 31*2(\bar{\pi})\}\}$ of descent views are the same. However, the descent views are easier to visualize for $\bar{\pi}$. In the middle diagram, they correspond to the rows containing the horizontal dashed lines, emanating towards the right from the dotted lines which indicate the descents.

The only descent view of π and $\bar{\pi}$ which is not a run bottom is 5. It will be convenient to encode the multiset of descent views which are not run bottoms as a sequence of multiplicities, which in the case at hand is $(d_\pi(v))_{v \in [n] \setminus RB} = (0, 1, 0, 0, 0)$.

Let us now consider the partial permutation $\sigma_e = 375??89??$, illustrated in the right-most diagram. An exceedance of σ_e corresponds to a column containing a dot strictly above the diagonal, and an inversion corresponds to a pair of dots such that the right dot is below the left dot. Thus, σ_e is only defined on the set of exceedances, which coincides with the set $[n] \setminus RT$.

More subtly, an invisible inversion corresponds to a pair of dots such that the right dot is below the left dot and the row containing the right dot intersects the column containing the left dot strictly above the diagonal. These are indicated by dashed horizontal lines emanating towards the right from the vertical dotted lines connecting the diagonal with the value at the exceedance. Thus, the only invisible inversion of σ_e is the pair $(2, 3)$, with $\sigma_e(3) = 5$. However, any permutation σ that agrees with σ_e on the set of exceedances and has no other exceedances will evidently have $\{\{2, 4, 4, 5, 6\}\}$ as a multiset of invisible inversion bottoms.

3. The exceedances of $\chi(\pi)$

We define the bijection $\chi : \mathfrak{S}_n \rightarrow \mathfrak{S}_n$ in two steps. First we describe the exceedances of $\sigma = \chi(\pi)$. The following is our key lemma.

Lemma 3.1. *Let RB and RT be the set of run bottoms, respectively run tops, of π . There exists a unique bijection $\sigma_e = \chi_e(\pi) : [n] \setminus RT \rightarrow [n] \setminus RB$ whose multiset of inversion bottoms is the multiset of descent views of π which are not run bottoms.*

This bijection satisfies $\sigma_e(i) > i$ for all $i \in [n] \setminus RT$.

Example 3.1. *Referring to the example illustrated in Figure 1, the lemma says that $375??89??$ is the unique partial permutation defined on the set $\{1, 2, 3, 6, 7\}$ with values in the set $\{3, 5, 7, 8, 9\}$ having 5 as the only inversion bottom. Moreover, this partial permutation is defined only on exceedances.*

Proof. Let $d_\pi(v)$ be the number of $\pi(i)$'s that make v a descent view, that is, the multiplicity of v in the multiset of descent views. Let $v \in [n] \setminus RB$. Thus, $d_\pi(v)$ is the number of positions i such that $\pi(i + 1) < v < \pi(i)$ and the run containing $\pi(i + 1)$ and $\pi(i)$ has a smaller run bottom than the run containing v . Therefore, if R is the set of runs of π and b is the run bottom of the run containing v ,

$$d_\pi(v) = \#\{r \in R \mid \min(r) < v < \max(r) \text{ and } \min(r) < b\}.$$

We now define σ_e iteratively. Let $S = [n] \setminus RT$. For $v \in [n] \setminus RB$, beginning with the smallest, we let $\sigma_e^{-1}(v)$ be the $(d_\pi(v) + 1)$ -st element in S and then remove this element from S .

When defining $\sigma_e^{-1}(v)$, there are $|[n] \setminus RB| - \#\{w \in [n] \setminus RB \mid w < v\}$ elements in S , which is at least $d_\pi(v) + 1$ because

$$d_\pi(v) + 1 \leq \#\{w \in [n] \setminus RB \mid w > v\} + 1 = |[n] \setminus RB| - \#\{w \in [n] \setminus RB \mid w < v\},$$

which ensures that σ_e is well-defined.

Uniqueness follows from the fact that v is the inversion bottom of precisely those $d_\pi(v)$ elements whose preimage is in what remains of S but is smaller than $\sigma_e^{-1}(v)$.

It remains to show that $\sigma_e^{-1}(v) < v$. Initially, there are

$$v - 1 - \#\{w \in RT \mid w < v\}$$

elements in S that are strictly less than v . When defining $\sigma_e^{-1}(v)$, we can assume by induction that only elements strictly less than v were removed from S . Moreover, the number of removed elements equals $\#\{w \in [n] \setminus RB \mid w < v\}$. Thus, the number of elements strictly less than v in what remains of S is

$$v - 1 - \#\{w \in RT \mid w < v\} - \#\{w \in [n] \setminus RB \mid w < v\}.$$

It remains to show that this is larger than $d_\pi(v)$, which follows from these two observations:

$$v - 1 - \#\{w \in [n] \setminus RB \mid w < v\} = \#\{w \in RB \mid w < v\},$$

and

$$\#\{w \in RB \mid w < v\} - \#\{w \in RT \mid w < v\} = \#\{r \in R \mid \min(r) < v \leq \max(r)\} > d_\pi(v).$$

□

Lemma 3.2. *Let π_1 and π_2 be two permutations. Then*

$$\chi_e(\pi_1) = \chi_e(\pi_2) \Leftrightarrow R(\pi_1) = R(\pi_2).$$

Example 3.2. *Referring to the example illustrated in Figure 1, there are twelve permutations with the same set of runs as π , and twelve permutations whose restriction to $[n] \setminus RT$ equals $\chi_e(\pi)$ and which do not have any further exceedances. This is consistent with the lemma.*

Proof. If the sets of runs of π_1 and π_2 coincide we have $\chi_e(\pi_1) = \chi_e(\pi_2)$, because the multiset of descent views of a permutation only depends on its set of runs.

Conversely, $\chi_e(\pi_1) = \chi_e(\pi_2)$ immediately implies that the sets of run bottoms and run tops of π_1 and π_2 are the same. We show how to reconstruct the set of runs of a permutation π from $\chi_e(\pi)$. Let us first determine the matching between run bottoms and run tops.

Let t be a run top of π , beginning with the smallest. Suppose that t is not a run bottom. In this case $d_\pi(t)$ is the multiplicity of t in the multiset of inversion bottoms of $\chi_e(\pi)$. Then, as t completes $d_\pi(t)$ descent views, t must be the run top for the $(d_\pi(t) + 1)$ -st smallest run bottom which is not yet matched with a run top. If t is a run bottom, $d_{\pi_1}(t)$ is not determined by $\chi_e(\pi_1)$. However, in this case it is clearly a singleton run in both π_1 and π_2 .

Finally, suppose the v is neither a run bottom nor a run top. Then, among the runs with smaller run bottom and larger run top, v belongs to the one with $(d_\pi(v) + 1)$ -st smallest run bottom. □

Lemma 3.3. *Let π be a permutation and let σ be any permutation that agrees with $\chi_e(\pi)$ on the set of exceedances and has no other exceedances. Then the multiset of invisible inversion bottoms of σ equals the multiset of descent views of π .*

Example 3.3. *Referring to the example illustrated in Figure 1, we see that the multisets of rows containing dashed lines is the same in the middle diagram and the right-most diagram, in agreement with the lemma.*

Proof. Let RB be the set of run bottoms of π . If $v \in [n] \setminus RB$, then by Lemma 3.1 v is the value of an exceedance of σ . Therefore, v is an invisible inversion bottom if and only if it is an inversion bottom. Thus, the multiplicity of v in the multiset of invisible inversion bottoms equals the multiplicity of v in the multiset of descent views of π .

Suppose now that $v \in RB$. We partition the set $[v - 1]$ in two ways. For any element w let $rt(w)$ be the run top in the decreasing run containing w . Then the first partition of $[v - 1]$ is as follows:

$$\begin{aligned} A &= \{w \in [v - 1] \mid w \notin RB\} \\ B &= \{w \in [v - 1] \mid w \in RB \text{ and } rt(w) \leq v\} \\ C &= \{w \in [v - 1] \mid w \in RB \text{ and } rt(w) > v\}. \end{aligned}$$

Note that $\#C = d_\pi(v)$. For the second partition of $[v - 1]$, let

$$\begin{aligned} F &= \{i \in [v - 1] \mid \sigma(i) \leq i\} \\ G &= \{i \in [v - 1] \mid \sigma(i) > i \text{ and } \sigma(i) \leq v\} \\ H &= \{i \in [v - 1] \mid \sigma(i) > i \text{ and } \sigma(i) > v\}. \end{aligned}$$

Note that $\#H$ is the multiplicity of v in the multiset of invisible inversions bottoms, because with $j = \sigma^{-1}(v) \geq v$ we have $i < j$ and $i < v = \sigma(j) < \sigma(i)$ for any $i \in H$. Moreover, $\#F = \#B$, since every element i of F is a run top of π , and mapping i to the corresponding run bottom yields a bijection. Finally, $\#G = \#A$, with σ restricted to G serving as a bijection: because any $i \in G$ is an exceedance of σ we find that $\sigma(i)$ is not a run bottom of π and therefore $\sigma(i) \neq v$.

We conclude that $\#C = \#H$, which is what we wanted to prove. □

Proposition 3.1. *Let π be a permutation and let σ be any permutation that agrees with $\chi_e(\pi)$ on the set of exceedances and that has no other exceedances. Then the sets of global ascents of π and σ are the same.*

Example 3.4. *Referring to the example illustrated in Figure 1, π evidently has no global ascents, and it is easy to check visually that the same is true for any permutation extending σ_e .*

Proof. If π has a global ascent at index m , $\bar{\pi} = (\pi(1), \dots, \pi(m))$ is a permutation of $[m]$. Moreover, any pair $(i, j) \in 31*2(\pi)$ has either both i and j in $[m]$, or both in the complement of $[m]$. The same is true for pairs $(i, j) \in \text{Inv}(\pi)$.

Therefore, the restriction of $\chi_e(\pi)$ to $[m]$ is the same as $\chi_e(\bar{\pi})$. Since the remaining values of σ are weak deficiencies, σ also has a global ascent at index m .

Conversely, suppose that σ has a global ascent at index m . Since σ agrees with $\chi_e(\pi)$ on the set of exceedances and has no other exceedances, it maps run tops of π to run bottoms of π . Because it has a global ascent at m , any run top larger than m is mapped to a run bottom larger than m . Therefore, the number of run tops of π smaller than m must be equal to the number of run bottoms of π smaller than m . Thus, any run of π has either both its run bottom and its run top and therefore all of its elements contained in $[m]$ or all of its elements contained in the complement of $[m]$. Finally, any ordering of the runs of π must be such that the runs contained in $[m]$ must come first. This in turn implies that π has a global ascent. \square

4. The weak deficiencies of $\chi(\pi)$

Lemma 3.1 shows that the values of σ on $[n] \setminus RT$ are determined by the first three requirements in Theorem 1.1. This is not the case for the values on the run tops of π , even when imposing the final requirement stated in Theorem 1.1.

In this section, we provide two different ways to complete the definition of σ . Note that, if we ignore the final requirement from Theorem 1.1, it suffices to encode the order of the runs of π in the bijection $\sigma_d : RT \rightarrow RB$.

In both variants we will make use of the following construction. Let f be a partial permutation and let $c \in [n]$ be an element not contained in a cycle of f . Then the *iterated preimage* $f^*(c)$ of c is the element $s \in S$ such that $f^k(s) = c$, with $k \geq 0$ maximal. In particular, if c is not in the image of f , then $f^*(c) = c$.

We define the values of σ on the set of run tops iteratively, beginning with the smallest run top and removing the corresponding runs from π . Thus, throughout the algorithm, we consider σ as a partial permutation, which initially equals σ_e . Note that σ_e does not contain any cycles.

The first three steps of both variants are identical:

- (1) Let t be the smallest run top in π .
- (2) Remove the run containing t from π .
- (3) If t was the first element of π , set $\sigma(t) = \sigma^*(t)$.

In the following two sections, we provide two different ways of specifying $\sigma(t)$ if there is a run bottom to the left of t . In both cases, we set $\sigma(t) = \sigma^*(t')$ for some run top $t' > t$ satisfying $\sigma^*(t') < t$. Thus, throughout the algorithm, the functional digraph of σ is a collection of paths beginning with a run bottom and ending with a run top, which is also the largest element of the path, together with a collection of cycles. In each iteration, a run top is connected to a smaller run bottom. Assuming this, we can already conclude that the bijection satisfies the final requirement of Theorem 1.1.

Lemma 4.1. *A cycle of σ is completed in an iteration of the algorithm if and only if the smallest run top t is the first element of π . In this case, t is the maximal element of the cycle.*

Therefore, the left-to-right maxima of π are the maximal elements of the cycles of σ .

Proof. If the minimal run top t is the first element of π , and therefore a left-to-right maximum, we complete a cycle of σ by setting $\sigma(t) = \sigma^*(t)$. This cycle has maximal element t because only run tops $t' < t$ have been assigned an image yet, and, by assumption, all run tops are weak deficiencies.

Otherwise, we set $\sigma(t) = \sigma^*(t')$ for some run top $t' > t$. Since the (partial) map σ is injective, $\sigma^*(t')$ must be different from $\sigma^*(t)$ before this assignment, which means that no cycle is completed. \square

4.1 First variant

In the following, we will use an auxiliary map τ , which we now define. Let $\rho : RB \rightarrow RT$ be the bijection that maps each run bottom of π to the corresponding run top in the same run. Then, given σ , we set $\tau = \sigma^* \circ \rho$, restricted to the set of run bottoms remaining in π . Note that τ has to be recomputed whenever another value of σ is determined.

For example, if $\pi = 529687314$, we have

$$\rho = \begin{matrix} 1 & 2 & 4 & 6 \\ 8 & 5 & 4 & 9 \end{matrix} \text{ and } \sigma_e^* = \begin{matrix} 8 & 5 & 4 & 9 \\ 6 & 1 & 4 & 2 \end{matrix}$$

and therefore

$$\tau = \begin{matrix} 1 & 2 & 4 & 6 \\ 6 & 1 & 4 & 2 \end{matrix}.$$

We now specify the final step of the algorithm as follows:

(4.1) Otherwise, let a be the run bottom left of t , and let $k \geq 1$ be minimal such that $\tau^k(a) < t$. Set $\sigma(t) = \tau^k(a)$.

Let us illustrate this procedure by resuming our example. Initially, we have

$$\pi = 529687314 \quad \sigma = \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 5 & 2 & 1 & 8 & 9 & & \end{matrix} \quad \tau = \begin{matrix} 1 & 2 & 4 & 6 \\ 6 & 1 & 4 & 2 \end{matrix}.$$

The smallest run top is $t = 4$. There is a run to the left of t , with run bottom $a = 1$. Since $\tau(a) = 6 > 4$, and $\tau^2(a) = 2 < 4$, we set $\sigma(4) = 2$.

Removing the run 4 from π and taking into account that now $\tau(6) = \sigma^*(9) = 4$ we obtain

$$\pi = 52968731 \quad \sigma = \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 5 & 2 & 1 & 8 & 9 & & \end{matrix} \quad \tau = \begin{matrix} 1 & 2 & 6 \\ 6 & 1 & 4 \end{matrix}.$$

The smallest run top is now $t = 5$. Since t is the first element of π , we set $\sigma(5) = \sigma^*(5) = 1$. We remove the run 52 from π and obtain

$$\pi = 968731 \quad \sigma = \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 5 & 2 & 1 & 8 & 9 & & \end{matrix} \quad \tau = \begin{matrix} 1 & 6 \\ 6 & 4 \end{matrix}.$$

The smallest run top is now $t = 8$. There is a run to the left of t , with run bottom $a = 6$. Since $\tau(a) = 4 < 8$, we set $\sigma(8) = 4$. We also remove the run 8731 from π , which leaves us with

$$\pi = 96 \quad \sigma = \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 5 & 2 & 1 & 8 & 9 & 4 & \end{matrix} \quad \tau = \begin{matrix} 6 \\ 6 \end{matrix},$$

and it remains to set $\sigma(9) = \sigma^*(9) = 6$.

Lemma 4.2. *At any stage of the algorithm and for any run bottom b we either have $\tau(b) = \sigma_e^*(\rho(b))$ or $\tau(b)$ is less than or equal to the minimal remaining run top of π .*

Proof. Suppose that the minimal run top is t and the corresponding run bottom is b .

If t is the first element of π , then, after setting $\sigma(t)$, only b is removed from the domain of τ , and no other values of τ are modified.

Otherwise, let $c = \tau^{k-1}(a)$, and let $d = \tau(c) = \sigma^*(\rho(c))$. Then after setting $\sigma(t) = d$, the value of $\tau(c)$ becomes

$$\sigma^*(\rho(c)) = \sigma^*(\sigma^{-1}(d)) = \sigma^*(t) = \sigma^*(\rho(b)) = \tau(b).$$

If $\tau(b)$ still has its original value, we have

$$\tau(b) = \sigma_e^*(\rho(b)) = \sigma_e^*(t) \leq t.$$

Otherwise, $\tau(b)$ is not larger than the minimal run top remaining when it was modified itself, which in turn must be smaller than t . \square

Lemma 4.3. *In step (4.1), there exists an integer $k \geq 1$ such that $\tau^k(a) < t$. We then also have that $\tau^k(a) = \sigma^*(t')$ for some run top $t' > t$.*

Thus, the algorithm is well defined and satisfies the assumption.

Proof. Initially, τ is a permutation of the run bottoms of π . Thus, if $\tau^\ell(a) = (\sigma_e^* \circ \rho)^\ell(a)$ for all ℓ , the cycle of τ beginning with $\tau(a)$ also contains a , which is strictly smaller than t .

Otherwise, if $\tau^\ell(a) \neq (\sigma_e^* \circ \rho)^\ell(a)$ for some ℓ , then $\tau^\ell(a) \leq s < t$ by Lemma 4.2, where s is the minimal run top of what remained of π when the modification of $\sigma^*(\rho(\tau^{\ell-1}(a)))$ occurred.

It remains to show that $\tau^k(a) = \sigma^*(t')$ for some run top $t' > t$. To this end, let $t' = \sigma^m(\tau^k(a))$ with m maximal. Then $t' \geq t$, because all run tops strictly smaller than t have already been assigned an image.

Suppose that $t' = t$, that is, $\tau^k(a) = \sigma^*(t) = \sigma^*(t)$. Then $\tau^{k-1}(a) = \rho^{-1}(t) \leq t$. By the minimality of k we obtain that $\tau^{k-1}(a) = t = \rho(t)$. However, if $\tau(c) = \sigma^*(\rho(c)) = t$ then $\rho(c) = t$, because t has no image, which in turn implies that t must be a fixed point of τ . This contradicts the assumption that $\tau^k(a) = \sigma^*(t) = t$. \square

Proposition 4.1. *χ is injective, and therefore a bijection.*

Proof. Suppose that $\chi(\pi_1) = \chi(\pi_2)$. By Lemma 3.2, π_1 and π_2 have the same set of runs, which implies that the same τ will be used for both to compute the weak deficiencies of $\chi(\pi_1)$ and $\chi(\pi_2)$. We show that the sequence of runs is the same in π_1 and π_2 .

Suppose that there is a run, with run top t and run bottom b , which is the first run in what remains of π_1 but in what remains of π_2 the run bottom a precedes t . Then $\chi(\pi_1)(t) = \sigma^*(t) = \tau(b) = \tau^k(a) = \chi(\pi_2)(t)$. Since τ is injective, we therefore have $\tau^{k-1}(a) = b \leq t$, which contradicts the minimality of k .

Finally, suppose that the currently minimal run top t is preceded by a_1 in what remains of π_1 , and by a_2 in what remains of π_2 . Then we have $\chi(\pi_1)(t) = \tau^{k_1}(a_1) = \tau^{k_2}(a_2) = \chi(\pi_2)(t)$. Without loss of generality, we may assume that $k_1 < k_2$. Since τ is injective, $\tau^{k_2-k_1}(a_2) = a_1 \leq t$, contradicting the minimality of k_2 . \square

4.2 Second variant

We specify the final step of the algorithm as follows:

- (4.2) Otherwise, let i be the index of the run bottom left of t in the sequence $b_1 < b_2 < \dots$ of run bottoms in what remains of π and let s_i be the i -th smallest element other than $s = \sigma^*(t)$ which is not yet in the image of σ . Set $\sigma(t) = s_i$.

Again, let us illustrate this procedure by resuming our example. Initially, we have

$$\pi = 52\,96\,8731\,4 \quad \sigma = \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 5 & 1 & 4 & 8 & 9 & & \end{matrix}$$

The smallest run top is $t = 4$, with corresponding run bottom $b = 4$. There is a run to the left of t , whose run bottom 1 is the smallest among the run bottoms in π other than $b = 4$. The first element other than $\sigma^*(t) = 4$ which is not in the image of σ is 1, so we set $\sigma(4) = 1$.

Removing the run 4 from π we obtain

$$\pi = 52\,96\,8731 \quad \sigma = \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 5 & 1 & 4 & 8 & 9 & & \end{matrix}$$

The smallest run top is now $t = 5$. Since t is the first element of π , we set $\sigma(5) = \sigma^*(5) = 4$.

We remove the run 52 from π , to obtain

$$\pi = 96\,8731 \quad \sigma = \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 5 & 1 & 4 & 8 & 9 & & \end{matrix}$$

The smallest run top is now $t = 8$, with corresponding run bottom $b = 1$. There is a run to the left of t , whose run bottom 6 is the smallest among the run bottoms in π other than $b = 1$. The first element other than $\sigma^*(t) = 6$ which is not in the image of σ is 2, so we set $\sigma(8) = 2$.

Finally, we remove the run 8731 from π . This leaves us with

$$\pi = 96 \quad \sigma = \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 3 & 7 & 5 & 1 & 4 & 8 & 9 & 2 & \end{matrix}$$

and it remains to set $\sigma(9) = \sigma^*(9) = 6$.

Lemma 4.4. *In step (4.2) of the algorithm, the number of run bottoms strictly less than t equals the number of elements strictly less than t other than s which are not in the image of σ .*

Thus, the algorithm is well defined. Moreover, the assumption is satisfied.

Proof. We consider the effect of a single iteration of the algorithm on the set of run bottoms of what remains of π and the set of elements that are not yet in the image of σ . To do so, initialize the two sets B and S with the initial set of run bottoms RB of π . After each iteration, remove the run bottom b in the same run as t from B and remove $\sigma(t)$ from S . Thus, only elements smaller than or equal to t are removed from B .

If step (4) is executed, also an element smaller than or equal to t is removed from S , because $\sigma^*(t) \leq t$. On the other hand, suppose that in previous iterations only elements smaller than or equal to the then-current run top were removed from S . Then, in step (4), the number of elements in $S \setminus s$ which are less than or equal to t equals the number of elements in $B \setminus b$ which are less than or equal to t . Since $b_i \in B \setminus b$, we conclude that $s_i \in S \setminus s$.

Finally, since $s_i \neq \sigma^*(t)$ and all elements smaller than t have an image, $s_i = \sigma(t')$ for some run top strictly larger than t . □

Proposition 4.2. *χ is injective, and therefore a bijection.*

Proof. Suppose that $\chi(\pi_1) = \chi(\pi_2)$. By Lemma 3.2, π_1 and π_2 have the same set of runs. We show that the sequence of runs is the same in π_1 and π_2 .

Suppose that there is a run, with run top t , which is the first run in what remains of π_1 but in what remains of π_2 the run bottom b_i precedes t . Then $s = \sigma^*(t) = s_i$, which is impossible.

Otherwise, suppose that the minimal run top t is preceded by b_i in what remains of π_1 , and by b_j in what remains of π_2 . Then we have $s_i = s_j$ and therefore also $i = j$. □

References

- [1] P. Alexandersson and O. Nabawanda, *Peaks are preserved under run-sorting*, In: Enumer. Comb. Appl. 2.1 (2022), Paper No. S2R2.
- [2] R. J. Clarke, E. Steingrímsson, and J. Zeng, *New Euler-Mahonian statistics on permutations and words*, Adv. in Appl. Math. 18.3 (1997), 237–270.
- [3] M. Coopman and Z. Hamaker, *Involutions under Bruhat order and labeled Motzkin paths*, European J. Combin. 103 (2022), 103513.
- [4] Z. Hamaker, E. Marberg, and B. Pawłowski, *Fixed-point-free involutions and Schur P -positivity*, J. Comb. 11:1 (2020), 65–110.