FPSAC 2012, Nagoya, Japan DMTCS proc. **AR**, 2012, 45–56

Modified Growth Diagrams, Permutation Pivots, and the BWX Map ϕ^*

Jonathan Bloom[1] and Dan Saracino[2]

[1] *Dartmouth College, Hanover, NH, USA*
[2] *Colgate University, Hamilton, NY, USA*

Abstract. In their paper on Wilf-equivalence for singleton classes, Backelin, West, and Xin introduced a transformation ϕ^*, defined by an iterative process and operating on (all) full rook placements on Ferrers boards. Bousquet-Mélou and Steingrímsson proved the analogue of the main result of Backelin, West, and Xin in the context of involutions, and in so doing they needed to prove that ϕ^* commutes with the operation of taking inverses. The proof of this commutation result was long and difficult, and Bousquet-Mélou and Steingrímsson asked if ϕ^* might be reformulated in such a way as to make this result obvious. In the present paper we provide such a reformulation of ϕ^*, by modifying the growth diagram algorithm of Fomin. This also answers a question of Krattenthaler, who noted that a bijection defined by the unmodified Fomin algorithm obviously commutes with inverses, and asked what the connection is between this bijection and ϕ^*.

Résumé. Dans leur article sur l'équivalence de Wilf pour les classes de singletons, Backelin, West et Xin ont introduit une transformation ϕ^*, définie par un processus itératif et opérant sur (tous) les placements complets de tours sur un plateau de Ferrers. Bousquet-Melou et Steíngrimsson ont démontré l'analogue du résultat principal de Backelin, West et Xin dans le contexte d'involutions, et pour ce faire ont eu besoin de démontrer que ϕ^* commute avec l'opération inverse. La preuve de cette commutativité est longue et difficile, et Bousquet-Melou et Steingrömsson se demandèrent s'il n'était pas possible de reformuler ϕ^* de sorte que le resultat devienne évident. Dans le présent article, nous proposons une telle reformulation de ϕ^*, en modifiant l'algorithme de croissance de diagramme de Fomin. Cette reformulation répond également à une question de Krattenthaler, qui, remarquant qu'une bijection définie par l'algorithme de Fomin non modifié commute évidemment avec l'opération inverse, se demanda quel était le rapport entre cette bijection et ϕ^*.

Keywords: Wilf-equivalence, RSK correspondence, Growth Diagrams, Bijection, Permutations

1 Introduction

For any permutation $\tau = \tau_1\tau_2 \dots \tau_r$, let $S_n(\tau)$ denote the set of permutations in S_n that avoid τ, in the sense that they have no subsequence order-isomorphic to τ.

In their paper *Wilf-equivalence for singleton classes* [1], Backelin, West, and Xin prove an important general result about permutations avoiding a single pattern: If $k, \ell \geq 1$ and ρ is a permutation of $\{k + 1, \dots, k + \ell\}$, then for every $n \geq k + \ell$ we have $|S_n(12 \dots k\rho)| = |S_n(k \dots 1\rho)|$. The key tool in the proof is a map ϕ^*, which operates on a permutation σ as follows: Order the $k \dots 1$-patterns $\sigma_{i_1} \dots \sigma_{i_k}$ in σ

lexicographically (according to the σ_j's, not the j's) and let $\phi(\sigma)$ be obtained from σ by taking the smallest $\sigma_{i_1} \dots \sigma_{i_k}$ and placing in positions $i_1, \dots, i_k$ the values $\sigma_{i_2}, \dots, \sigma_{i_k}, \sigma_{i_1}$, respectively, and leaving all other entries of σ fixed. Let $\phi^*(\sigma)$ be obtained by applying ϕ repeatedly until no $k \dots 1$-patterns remain. It is shown that ϕ^* induces a bijection from $S_n(k-1 \dots 1k)$ onto $S_n(k \dots 1)$, and that from this bijection one can obtain an important ingredient of the proof, namely a bijection from $S_n(1 \dots k)$ onto $S_n(k \dots 1)$.

In [4], Bousquet-Mélou and Steingrímsson prove the analogue of the Backelin-West-Xin Theorem in the context of involutions, and in so doing they must prove that $\phi^*(\sigma^{-1}) = (\phi^*(\sigma))^{-1}$ for all permutations σ, so that the bijection from $S_n(1 \dots k)$ onto $S_n(k \dots 1)$ will also commute with inverses. The proof of the commutation result for ϕ^* is long and difficult, and Bousquet-Mélou and Steingrímsson ask for a reformulation of ϕ^* that will make this result obvious. In [7], Krattenthaler describes another bijection from $S_n(1 \dots k)$ into $S_n(k \dots 1)$, in terms of growth diagrams, and notes that this bijection clearly commutes with inverses. He asks what connection there is between this bijection and ϕ^*. In the present paper we answer both questions, by providing a reformulation of ϕ^* in terms of growth diagrams.

In proving their theorem, Backelin, West, and Xin find it necessary to work in the context of full rook placements on Ferrers boards, which includes permutations as a special case. For any Ferrers board F and any permutation τ, let $S_F(\tau)$ denote the set of all full rook placements on F that avoid τ. (The relevant definitions will be reviewed in Section 2.)

Backelin, West, and Xin prove that $|S_F(1 \dots k\rho)| = |S_F(k \dots 1\rho)|$, for all F and all permutations ρ of $\{k+1, \dots, k+\ell\}$, and in so doing they use an extension of ϕ^* to full rook placements. Bousquet-Mélou and Steingrímsson also use this extension, so they prove that ϕ^* commutes with inverses in this broader context. Accordingly, our reformulation of ϕ^* will be given in this context, or rather in the even broader context of arbitrary rook placements (not necessarily full), with the term "inverse" interpreted appropriately.

The outline of this extended abstract is as follows. In Section 2 we review the needed background material on the Robinson-Schensted correspondence for partial permutations, Ferrers boards and rook placements, and growth diagrams. In Section 3 we give our reformulation of ϕ^* and the proof that it works, modulo a "Main Lemma". In Section 4 we introduce the tool that will be used to prove this lemma: the "pivots" of a rook placement on a rectangular Ferrers board. The pivots are related to the "L-corners" of [2], and are a generalization of the "rcL-corners" of [9]. (We have chosen to use the term "pivots", instead of "corners", because of the prior use of the term "corners" in connection with the diagram of a permutation.) Section 5 contains the proof of the Main Lemma. Although omitted in the abstract our paper [3] contains a concluding Section 6 indicates how the proof leads naturally to a notion of generalized Knuth transformations.

2 Background

2.1 Ferrers boards, rook placements and ϕ^*

Consider an $n \times n$ array of squares, and identify the pair (i, j) with the square located in the ith column from the left and the jth row from the bottom. For any square (i, j) in the array, let $R(i, j)$ denote the rectangle consisting of all squares (k, ℓ) such that $k \le i$ and $\ell \le j$. A *Ferrers board* (in French notation) is any subset F of such an array with the property that for all $(i, j) \in F$ we have $R(i, j) \subseteq F$. So for some t and some $\lambda_1 \ge \dots \ge \lambda_t$, the Ferrers board consists of the first λ_j squares from the jth row of the array,

$1 \leq j \leq t$. The *conjugate* of F is the Ferrers board $F' = \{(j,i) : (i,j) \in F\}$, so that F' is obtained by reflecting F across the SW-NE diagonal.

A *rook placement* on a Ferrers board F is a subset of F that contains at most one square from each row of F and at most one square from each column of F. We indicate the squares in the placement by putting markers (e.g., dots or X's) in them. A rook placement is called *full* if it includes exactly one square from each row and column of F. (So if there exists a full rook placement on F, then F has the same number of columns as rows.) From any rook placement P on F there results a partial permutation π such that square (i,j) is in P if and only if j is the value of the bijection π at input i. P is a full placement if and only if F has n rows and n columns for some n and π is a permutation of $\{1,\ldots,n\}$. For any rook placement P on F, the *inverse* P' of P is the placement on the conjugate board F' obtained by reflecting F and all the markers for P across the SW-NE diagonal. The partial permutations resulting from P and P' are inverses of each other, if we regard them as bijections between sets. If they are permutations, they are inverses in the usual sense.

We say that a rook placement P *contains* a permutation $\tau \in S_r$ if and only if the resulting partial permutation π contains a subsequence $\pi_{i_1} \ldots \pi_{i_r}$ order isomorphic to τ such that there is a rectangular subboard of F that contains all the squares (i_j, π_{i_j}). In this case we refer to the sequence of squares (i_j, π_{i_j}) as an *occurrence* of τ in P. We say that P *avoids* τ if P does not contain τ.

It is clear how to extend the definitions of ϕ and ϕ^* to rook placements, by using only the occurrences of $k \ldots 1$ in P, in the sense of the preceding paragraph.

2.2 Growth diagrams

Our reformulation of ϕ^* will be accomplished by modifying Fomin's ([5,6], see also [7]) construction of the growth diagram of a rook placement P on a Ferrers board F.

Fomin's construction assigns partitions to the corners of all the squares in F, using the markers of P, in such a way that the partition assigned to any corner either equals the partition to its left or is obtained from it by adding 1 to one entry, and the partition assigned to any corner either equals the partition below it or is obtained from this partition by adding 1 to one entry. We start by assigning the empty partition $\emptyset$ to each corner on the left and bottom edges of F. We then assign partitions to the other corners inductively. Assuming that the northwest, southwest, and southeast corners of a square (i,j) have been assigned partitions NW, SW, and SE, we assign to the northeast corner the partition NE determined by the following rules.

1. If $NW \neq SE$ then let $NE = NW \cup SE$, the partition whose ith entry is the maximum of the ith entries of NW and SE. (Here we regard the absence of an entry as the presence of an entry 0.)
2. If $SW \neq NW = SE$ then NW is obtained from SW by adding 1 to the ith entry of SW, for some i. We obtain NE from NW by adding 1 to the $(i+1)$th entry.
3. If $SW = NW = SE$ then we let $NE = SW$ unless the square (i,j) contains a marker, in which case we obtain NE from SW by adding 1 to the first entry.

For an example of these rules see Figure 1.

Lemma 1 *([7, Theorem 5.2.4] or [12, Theorem 7.13.5]). For any square (i,j), the partition assigned to the northeast corner of (i,j) is the shape of the Robinson-Schensted tableaux for the partial permutation resulting from the restriction of P to the rectangle $R(i,j)$.*

In general, if we are given F and the partitions in the growth diagram for P that occur along the right/up border of F (i.e., the border of F minus the horizontal bottom edge and the vertical left edge), we can inductively reconstruct the rest of the growth diagram and the placement P.

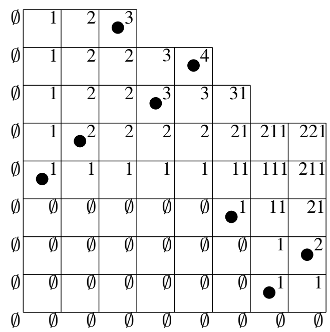

Fig. 1: An example of Fomin's growth diagram algorithm (GDA).

3 The Reformulation of ϕ^*

We now modify the growth diagram algorithm (GDA) of subsection 2.3 to get a new algorithm GDA_k for any $k \geq 2$. GDA_k retains rules (1) and (3) of GDA, but replaces rule (2) by the following variant.

2_k. Apply rule (2) with the proviso that if rule (2) produces a NE with k (nonzero) entries then delete the last entry and increase the first entry by 1.

The motivation for rule (2_k) comes from the theorem of Schensted which relates the number of parts of the insertion tableau with the length of the longest decreasing sequence. Therefore keeping the number of entries in a partition λ less than k prevents decreasing subsequences of length k in partial permutations whose Robinsion-Schensted tableaux have shape λ.

Definitions For any rook placement P on a Ferrers board F, let $\text{seq}(P, F)$ (respectively, $\text{seq}_k(P, F)$) denote the sequence of partitions along the right/up border of F that results from the application of GDA (respectively, GDA_k) to P and F.

Main Theorem Fix $k \geq 2$ in the definition of ϕ^*. Then for any rook placement P on a Ferrers board F,

$$\text{seq}_k(P, F) = \text{seq}(\phi^*(P), F).$$

Corollary For any rook placement P on a Ferrers board F,

$$\phi^*(P') = (\phi^*(P))'.$$

Proof. This is essentially clear from the fact that the algorithms GDA and GDA_k commute with the operation of taking the inverse of a placement.

In a bit more detail, $\mathrm{seq}_k(P', F')$ is the reverse of $\mathrm{seq}_k(P, F)$, so, by the Main Theorem, $\mathrm{seq}(\phi^*(P'), F')$ is the reverse of $\mathrm{seq}(\phi^*(P), F)$, and this reverse is $\mathrm{seq}((\phi^*(P))', F')$. By rules (A), (B), and (C) for the inverse algorithm for GDA, we conclude that $\phi^*(P') = (\phi^*(P))'$. $\square$

We will now give an example, to illustrate the Main Theorem and indicate the structure of its proof.

Example Consider the placement P (left) and $\phi^*(P)$ (right) with $k = 3$.

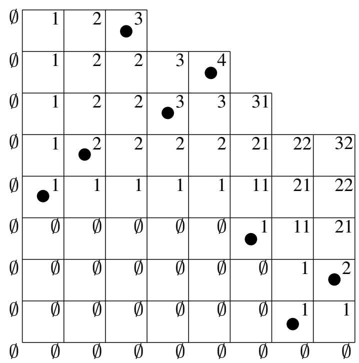

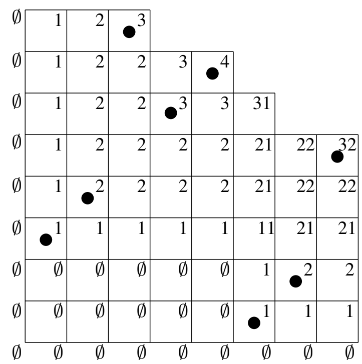

The partitions on the left are obtained by performing GDA_3 while the partitions on the right are obtained by performing GDA. Observe that the partitions along the right/up border of F are the same in both cases, so $\mathrm{seq}_3(P, F)$=$\mathrm{seq}(\phi^*(P), F)$, although the partitions in the interiors of the diagrams are not always the same.

Lastly, consider (below) the partitions obtained by performing GDA_3 on $\phi(P)$.

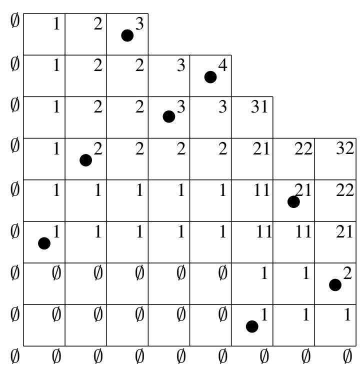

Notice that, while the results of performing GDA_3 on P and on $\phi(P)$ are not the same diagram, they do agree on the boundary of $R(7, 4)$, the smallest rectangular subboard of F that contains the 321-pattern on which ϕ acted and extends to the left and bottom edges of F. Because of the definition of GDA_3, this is enough to make the two diagrams agree everywhere outside the rectangle. The idea of the proof of the Main Theorem will be to show that performing GDA_k on P and on $\phi(P)$ yields the same partitions on the boundary of the smallest rectangular subboard of F that contains the $k \dots 1$-pattern on which ϕ acted and extends to the left and bottom edges of F.

Proof of the Main Theorem: We proceed by induction on the number of applications of ϕ required to compute $\phi^*(P)$. If no applications are required, the result is obvious, since $\phi^*(P) = P$ and performing GDA_k is the same as performing GDA.

Now suppose that m applications of ϕ are required to compute $\phi^*(P)$. Since computing $\phi^*(\phi(P))$ requires only $m-1$ applications of ϕ, we assume inductively that

$$\text{seq}_k(\phi(P), F) = \text{seq}(\phi^*(\phi(P)), F),$$

i.e., that

$$\text{seq}_k(\phi(P), F) = \text{seq}(\phi^*(P), F).$$

We want to show that $\text{seq}_k(\phi(P), F) = \text{seq}_k(P, F)$.

Let $R = R(a, b)$ be the smallest rectangular subboard of F that contains the $k \dots 1$-pattern on which ϕ acted to produce $\phi(P)$ and extends to the left and bottom edges of F. Let P_R and $\phi(P)_R$ be the restrictions of P and $\phi(P)$ to $R(a, b)$. By the definition of GDA_k, all we need to show is that

$$\text{seq}_k(\phi(P)_R, R) = \text{seq}_k(P_R, R).$$

To show this, we will use the following two lemmas, which are proved in [3].

Lemma 2 *The placement P_R on R contains no $k \dots 1$-pattern that begins in a row below the top row of R or ends in a column to the left of the rightmost column of R.*

Lemma 3 *The placement $\phi(P)_R$ on R contains no $k \dots 1$-pattern.*

By Lemma 3,

$$\text{seq}_k(\phi(P)_R, R) = \text{seq}(\phi(P)_R, R).$$

By Lemma 2, $\text{seq}_k(P_R, R) = \text{seq}(P_R, R)$ except possibly at the northeast corner of square (a, b).

Notation Let c_{ne}, c_{nw}, and c_{se} denote the northeast, northwest, and southeast corners of square (a, b).

To conclude the proof, it will suffice to prove the next lemma.

Main Lemma We have $\text{seq}(P_R, R) = \text{seq}(\phi(P)_R, R)$ except possibly at c_{ne}, the northeast corner of R.

For once this lemma is established, we will have

$$\text{seq}_k(\phi(P)_R, R) = \text{seq}_k(P_R, R)$$

except possibly at c_{ne}. To see that the two also agree at c_{ne}, let, by the Main Lemma, λ be the common value of $\text{seq}(P_R, R)$ and $\text{seq}(\phi(P)_R, R)$ at c_{nw}. Since

$$\text{seq}_k(\phi(P)_R, R) = \text{seq}(\phi(P)_R, R),$$

the value of $\text{seq}_k(\phi(P)_R, R)$ at c_{ne} is obtained from λ by adding 1 to the first entry. (This follows from the last statement in subsection 2.3 and the fact that $\phi(P)_R$ has a marker in square (a, b).) To see that $\text{seq}_k(P_R, R)$ has the same value at c_{ne}, let μ, ν be the values of $\text{seq}(P_R, R)$ at c_{se}, c_{ne}. Since, in the placement P_R on R, $R(a-1, b)$ and $R(a, b-1)$ contain no $k \ldots 1$-patterns (by Lemma 2) but $R(a, b)$ contains such a pattern, it must be that ν has k entries and each of λ, μ has $k-1$. Therefore, by the definition of GDA_k, the value of $\text{seq}_k(P_R, R)$ at c_{ne} is obtained by adding 1 to the first entry of λ.

The proof of the Main Lemma will occupy Section 5. Before continuing let us establish some useful notation for the sequel.

Notation For a square $B = (i, j)$ in a Ferrers board F, we denote i and j by $\text{col}(B)$ and $\text{row}(B)$, respectively.

4 Pivots

In this section we introduce the *left* and *right pivots* of rook placement on a rectangular Ferrers board, and show how they relate to the Robinson-Schensted correspondence and to ϕ^*. While the right pivots have nicer properties in connection with the RS correspondence, the left pivots have nicer properties with respect to ϕ^*.

Definitions Let P be a placement on rectangular Ferrers board F. We define the set of *left pivots of* P (respectively, *right pivots of* P) to be the placement $\text{piv}_l\ P$ (respectively, $\text{piv}_r\ P$) defined by inductively placing markers, row by row, from bottom to top, as follows.

First, there is no pivot in the bottom row. Now consider row $r > 1$. If there is no element of P in row r then we do not place a pivot in row r. Now suppose $X \in P$ is in row r. If there is a column to the left (respectively, right) of X that contains an element of P below row r but does not contain a pivot then we place a pivot in row r and in the rightmost (respectively, leftmost) such column.

Notation For a placement P on a rectangular Ferrers board F define $\text{rev}(P)$ to be the placement obtained by reflecting F and P along a vertical line.

Below is an example of left and right pivots of a placement.

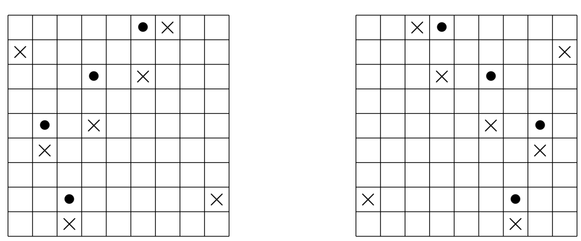

Fig. 2: On the left we have P and $\text{piv}_l(P)$. On the right we have $\text{rev}(P)$ and $\text{piv}_r(P)$. In both examples elements of P are denoted by $\times$ and the pivots are denoted by $\bullet$.

Looking at Figure 1 we see the following relationship between left and right pivots. Its proof is straightforward.

Lemma 4 *For any placement P on a rectangular Ferrers board F we have*

$$\mathrm{rev}\Big(\mathrm{piv}_l\ P\Big) = \mathrm{piv}_r\Big(\mathrm{rev}\ P\Big).$$

The utility of pivots is due to their connection with the Robinson-Schensted (RS) algorithm. Specifically, we prove that the RS algorithm applied to $\mathrm{piv}_r\ P$ gives the same insertion and recording tableaux, minus the top row, as the RS algorithm applied to P. And the insertion tableau of $\mathrm{piv}_l\ P$ is equal to $\mathrm{ins}\ P$, minus its left column. This is the content of the next two theorems.

Notation If Y is a standard Young tableau then we will denote by Y^- the tableau consisting of all but the top row of Y and ${}^-Y$ the tableau consisting of all but the left column of Y.

Theorem 1 *Let P be a placement on a rectangular Ferrers board. Then we have*

$$\mathrm{ins}(\mathrm{piv}_r\ P) = (\mathrm{ins}\ P)^- \qquad \textit{and} \qquad \mathrm{rec}(\mathrm{piv}_r\ P) = (\mathrm{rec}\ P)^-$$

For the proof see [3].

Remark 1 We thank Sergi Elizalde for suggesting to us that there might be a connection between our right pivots and Viennot's geometric construction. It turns out that the right pivots coincide with Viennot's "northeast corners". This follows from the fact that Viennot establishes the analogue of Theorem 1 for the northest corners.

To state our theorem relating left-pivots and the RS algorithm we first need the following.

Notation If Y is any standard Young tableau then denote by Y^{tr} the transposed tableau, i.e., the tableau obtained by reflecting Y across the NW-SE diagonal.

Likewise, if P is any placement on a rectangular Ferrers board F then denote by F^{tr} and P^{tr} the resulting board and placement obtained by reflecting F and P across the NW-SE diagonal.

Remark 2 Recall the well known fact that for any placement P we have $\mathrm{ins}(\mathrm{rev}\ P) = (\mathrm{ins}\ P)^{tr}$. For a proof of this result see [8], page 97.

Theorem 2 *Let P be a placement on a rectangular Ferrers board F. Then*

$$\mathrm{ins}(\mathrm{piv}_l\ P) = {}^-(\mathrm{ins}\ P).$$

Proof: By Lemma 4 we have $\mathrm{rev}(\mathrm{piv}_l\ P) = \mathrm{piv}_r(\mathrm{rev}\ P)$. Taking the insertion tableau of both sides and applying Theorem 1 to the right hand side we get

$$\mathrm{ins}\Big(\mathrm{rev}\ (\mathrm{piv}_l\ P)\Big) = [\mathrm{ins}\Big(\mathrm{rev}\ P\Big)]^-.$$

Then by Remark 2 we have $(\mathrm{ins}\ (\mathrm{piv}_l\ P))^{tr} = [(\mathrm{ins}\ P)^{tr}]^-$. By then transposing both sides we have

$$\mathrm{ins}(\mathrm{piv}_l\ P) = ([(\mathrm{ins}\ P)^{tr}]^-)^{tr} = {}^-(\mathrm{ins}\ P).$$

□

Convention Since we will only be working with left pivots for the remainder of this abstract, the term pivot will mean left pivot from now on.

The following definitions and lemmas will be useful in Sections 5. Their statements are given here but the proofs may be found in [3].

Definitions Let P be a placement on a rectangular Ferrers board. For any $X \in P$, if X is in the same column as some $V \in \mathrm{piv}_l\, P$ then define $\rho(P, X) = \mathrm{row}(V)$ else define $\rho(P, X) = \infty$.

Likewise, if X is in the same row as some $V \in \mathrm{piv}_l\, P$ then define $\kappa(P, X) = \mathrm{col}(V)$ else define $\kappa(P, X) = 0$.

If the placement is understood we will just write $\rho(X)$ or $\kappa(X)$.

Lemma 5 *If P is a placement on a rectangular Ferrers board F with n columns, then we have*

$$(\mathrm{piv}_l\, P)^{tr} = \mathrm{piv}_l(P^{tr})$$

Definition We say an increasing subsequence I of P is a *pivot-path*, if for all $i < |I|$, $\rho(I_i) = \mathrm{row}(I_{i+1})$, i.e., each consecutive pair creates a pivot.

Lemma 6 *Let S be a decreasing sequence of length k in P. Assume there is some $X \in P$ such that XS_1 is a 12-pattern. If $\rho(X) > \mathrm{row}(S_1)$ there exists a decreasing sequence D of length $k+1$ with $D_1 = X$, and if X and S_1 are the first and last elements of a pivot-path then there exists a decreasing sequence D of length k with $D_1 = X$.*

Note that the previous lemma and Lemma 5 directly imply the following result.

Lemma 7 *Let S be a decreasing sequence of length k in P. Assume there is some $X \in P$ such that S_1X is a 12-pattern. If $\kappa(X) < \mathrm{col}(S_1)$ then there exists a decreasing sequence D of length $k+1$ with $D_1 = X$.*

5 The Proof of the Main Lemma

The purpose of this section is to prove the Main Lemma which is the crucial piece needed to show that our growth diagram construction corresponds to the map ϕ^*. Note that up through Lemma 8, F will always denote a rectangular Ferrers board and P a rook placement on F.

We begin with some notation and definitions.

Notation Define $F|_{a,b}$ to be the Ferrers board consisting of all columns of F between and including columns a and b, and define $P|_{a,b}$ to be the restriction of P to $F|_{a,b}$.

Definitions Fix a subplacement S, let $k = |S|$, and order S such that $S = S_1 \dots S_k$ where $\mathrm{col}(S_i) < \mathrm{col}(S_{i+1})$. Further fix a column c and assume that P has no element in column c.

We first define a *left-shift*. Assume S lies entirely to the right of column c. Place markers in all the squares of P, and then shift each marker in square S_i for $1 < i \leq k$ horizontally left to column $\mathrm{col}(S_{i-1})$ and shift the marker in S_1 horizontally left to column c. Define $P(c \leftarrow S)$ to be the squares that now contain markers. Likewise, define $c \leftarrow S$ to be $P(S \leftarrow c) \setminus P$, the placement obtained from S by shifting it left.

Analogously, we define a *right-shift*. Assume S lies entirely to the left of column c. Place markers in all the squares of P, and then shift each marker in square S_i for $1 \leq i < k$ horizontally right to column $\mathrm{col}(S_{i+1})$ and shift the marker in S_k horizontally right to column c. Define $P(S \to c)$ to be the squares that now contain markers. Likewise, define $S \to c$ to be $P(S \to c) \setminus P$, the placement obtained from S by shifting it right.

Definitions Let P be a placement on a rectangular Ferrers board F. Let $a < b$ be columns of F.

Let k be the length of the longest decreasing sequence in $P|_{a,b}$. Define $\mathbf{d}_{a,b}(P)$ (respectively, $\mathbf{D}_{a,b}(P)$) to be the smallest (respectively, largest) decreasing sequence, under the lexicographical ordering, in $P|_{a,b}$ of length k.

Likewise, let m be the length of the longest increasing sequence in $P|_{a,b}$. Define $\mathbf{i}_{a,b}(P)$ (respectively, $\mathbf{I}_{a,b}(P)$) to be the smallest (respectively, largest) increasing sequence, under the lexicographical ordering, in $P|_{a,b}$ of length m.

We are now ready to state and prove Lemma 8. Although Lemma 8 is the key driver behind the Main Lemma, its statement is more general then what is needed to prove the Main Lemma. The reason for this is that the precise statement of Lemma 8 is exactly what is needed in Section 6.

Lemma 8 *Let P be a placement on a rectangular Ferrers board F and assume P has no marker in column a. Let $S = \mathbf{d}_{a,b}(P)$ and $P^* = P(a \leftarrow S)$. Then we have*

$$\mathrm{ins}(P) = \mathrm{ins}(P^*).$$

Proof: First let $S^* = a \leftarrow S$, i.e., the shifted sequence, and $k = |S|$. Define b' to be the column containing S_k and consider the truncated board and placements $G = F|_{1,b'}$, $Q = P|_{1,b'}$ and $Q^* = P^*|_{1,b'}$. Now in order to show $\mathrm{ins}(P) = \mathrm{ins}(P^*)$, it will suffice to show

$$\mathrm{ins}(Q) = \mathrm{ins}(Q^*). \tag{1}$$

To see this consider the GDA applied to P and P^* on the full board F. By [12], Theorem 7.13.5, knowing (1) is the same as knowing that the partitions along the line $x = b'$ in the GDA of P and in the GDA of P^* are identical. But the placements P and P^* are identical east of the line $x = b'$. Therefore if (1) holds we know that the partitions along the right border of F in the GDA of P and in the GDA of P^* are also identical. By [12], Theorem 7.13.5, this implies that $\mathrm{ins}(P) = \mathrm{ins}(P^*)$.

In order to show (1) it is sufficient to prove

$$\mathrm{piv}_l\, Q = \mathrm{piv}_l\, Q^*. \tag{2}$$

To see why note that (2) along with Theorem 2 implies that $^-[\mathrm{ins}(Q)] = {}^-[\mathrm{ins}(Q^*)]$, which forces $\mathrm{ins}(Q) = \mathrm{ins}(Q^*)$ as needed.

To establish (2) we will proceed inductively by assuming that the pivots (if any) below some row r are unchanged and then showing that the pivot (if any) on row r is unchanged. To do so we consider two cases.

`Case 1`: Row r contains an element S_i of S.

Let R be the rectangular region of F containing all squares west of col S_i and south of row S_i. By the nature of a left-shift, $Q|_R = Q^*|_R$. Further, by our induction hypothesis $\mathrm{piv}_l(Q|_R) = \mathrm{piv}_l(Q^*|_R)$. It then follows that the only way the pivot (if any) on row r could change is if $\mathrm{piv}_l(Q)$ contains a pivot in row r between columns a' and $\mathrm{col}(S_i)$, where $a' = a$, if $i = 1$, and $a' = \mathrm{col}(S_{i-1})$, if $i > 1$.

To show this is impossible assume, for a contradiction, that $\mathrm{piv}_l(Q)$ does contain a pivot in row r between columns a' and $\mathrm{col}(S_i)$. Then there must exist some $X \in Q$ that is directly below this pivot. Since XS_i is a pivot-path, Lemma 6, applied to X and $S_iS_{i+1}\ldots$, implies the existence of a decreasing sequence D of length $k - i + 1$ with first element X. But this is a contradiction since the decreasing sequence $S_1 \ldots S_{i-1}D$ is smaller than S.

`Case 2`: Row r contains $X \in Q \setminus S$.

By the maximality of the length of S, $S_{i-1}XS_i$ cannot be a 321-pattern for any i. Therefore it will suffice to assume that S_iX is a 12-pattern for some i and that i is chosen as small as possible. Then $\mathrm{col}(S_i) \leq \kappa(X, Q)$. If not then Lemma 7 would give rise to sequence of length $k + 1$ contradicting the maximality of the length of S. Therefore we must have an element $Y \in Q$ with $\mathrm{col}(S_i) \leq \mathrm{col}(Y) = \kappa(X, Q)$. Since $\mathrm{col}(S_i) \leq \mathrm{col}(Y) < b'$ then column $\mathrm{col}(Y)$ must contain an element of Q^* below row r. (Note that for the case where $Y \in S$, Y cannot be the last element of S.) This plus the induction hypothesis implies that $\mathrm{col}(Y) \leq \kappa(X, Q^*)$. Assume now, for a contradiction, that $\mathrm{col}(Y) < \kappa(X, Q^*)$ and let $Z \in Q^*$ be such that $\mathrm{col}(Y) < \mathrm{col}(Z) = \kappa(X, Q^*)$. Now, by similar reasoning, $\mathrm{col}(Z)$ must contain an element of Q below row r. But this plus the induction hypothesis implies that $\kappa(X, Q) = \mathrm{col}(Y) < \mathrm{col}(Z) \leq \kappa(X, Q)$, an obvious contradiction. □

To prove the Main Lemma it suffices, by [12], Theorem 7.13.5, to prove the following.

Lemma 9 *Let P be a placement on a not necessarily rectangular Ferrers board. Let S be the smallest $k \ldots 1$-pattern in P and let $b = \mathrm{row}(S_1)$ and $a = \mathrm{col}(S_k)$. If $R_1 = R(a, b-1)$ then*

$$\mathrm{ins}(P|_{R_1}) = \mathrm{ins}(\phi(P)|_{R_1}). \tag{3}$$

Likewise, if $R_2 = R(a-1, b)$ then

$$\mathrm{rec}(P|_{R_2}) = \mathrm{rec}(\phi(P)|_{R_2}). \tag{4}$$

Proof: Let $Q = P|_{R_1}$ and $T = \phi(P)|_{R_1}$. Observe that $S_2, \ldots, S_k = \mathbf{d}_{\mathrm{col}(S_1),a}(Q)$. Now Lemma 8 implies (3).

For (4) let $R = R(a, b)$ and $Q = P|_R$ and $T = \phi(P)|_R$. Observe that $T' = \phi(Q')$, where Q' and T' are the inverses of Q and T in the sense of Section 2 (see Figure 3). This follows from the fact that S is the decreasing sequence of maximal length in R, which implies that S is also the position smallest sequence of length k in R (i.e., S is the smallest if we order $k \ldots 1$-patterns lexicographically according to the positions of the entries, rather than the values of the entries), so that S' is the value smallest decreasing sequence of length k in R'.

Applying (3), we know that $\mathrm{ins}(Q'|_{R_2'}) = \mathrm{ins}(T'|_{R_2'})$. So by the theorem of Schützenberger we have $\mathrm{rec}(Q|_{R_2}) = \mathrm{rec}(T|_{R_2})$. This completes the proof. □

Acknowledgements

The first author would like to thank his thesis advisor, Sergi Elizalde, for suggesting the problem of reformulating ϕ^*.

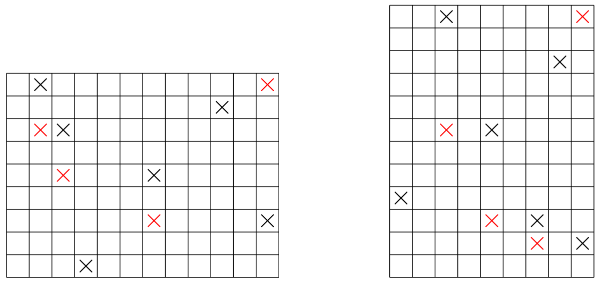

Fig. 3: On the left we have an example of Q and T. On the right we have Q' and T'. In both pictures the $\phi(S)$ is marked by the red $\times$'s.

References

[1] J. Backelin, J. West, G. Xin, Wilf-equivalence for singleton classes, *Adv. Appl. Math.* **38** (2007), 133–148.

[2] J. Bloom and D. Saracino, Another look at pattern-avoiding permutations, *Adv. Appl. Math.* **45** (2010), 395–409.

[3] J. Bloom and D. Saracino, Modified Growth Diagrams, Permutation Pivots, and the BWX map ϕ^*, `http://arxiv.org/abs/1103.0319`.

[4] M. Bousquet-Mélou, E. Steingrímsson, Decreasing subsequences in permutations and Wilf equivalence for involutions, *J. Algebraic Combin.* **22** (2005), 383–409.

[5] S. Fomin, Generalized Robinson-Schensted-Knuth correspondence, *Zapiski Nauchn. Sem. LOMI* **155** (1986), 156–175.

[6] S. Fomin, Schensted algorithms for dual graded graphs, *J. Alg. Comb.* **4** (1995), 5–45.

[7] C. Krattenthaler, Growth diagrams, and increasing and decreasing chains in fillings of Ferrers shapes, *Adv. Appl. Math.* **37** (2006), 404–431.

[8] B. Sagan, The Symmetric Group, second edition, Springer-Verlag, New York 2001.

[9] D. Saracino, On two bijections from $S_n(321)$ to $S_n(132)$, *Ars. Comb.* **101** (2011), 65-74.

[10] C. Schensted, Longest increasing and decreasing subsequences, *Canad. J. Math* **13** (1961), 179–191.

[11] M. Schützenberger, Quelques remarques sur une construction de Schensted, *Math. Scand.* **12** (1963), 117–128.

[12] R. Stanley, Enumerative Combinatorics, Vol. II, Cambridge Univ. Press, Cambridge, 1999.