EW-tableaux, Le-tableaux, tree-like tableaux and the Abelian sandpile model

Thomas Selig* Jason P. Smith† Einar Steingrímsson*

Department of Computer and Information Sciences,
University of Strathclyde,
Glasgow, G1 1XH, UK

{thomas.selig,jason.p.smith}@strath.ac.uk
einar@alum.mit.edu

Submitted: Nov 15, 2017; Accepted: Jun 28, 2018; Published: Jul 27, 2018

Abstract

An EW-tableau is a certain 0/1-filling of a Ferrers diagram, corresponding uniquely to an acyclic orientation, with a unique sink, of a certain bipartite graph called a Ferrers graph. We give a bijective proof of a result of Ehrenborg and van Willigenburg showing that EW-tableaux of a given shape are equinumerous with permutations with a given set of excedances. This leads to an explicit bijection between EW-tableaux and the much studied Le-tableaux, as well as the tree-like tableaux introduced by Aval, Boussicault and Nadeau.

We show that the set of EW-tableaux on a given Ferrers diagram are in 1-1 correspondence with the minimal recurrent configurations of the Abelian sandpile model on the corresponding Ferrers graph.

Another bijection between EW-tableaux and tree-like tableaux, via spanning trees on the corresponding Ferrers graphs, connects the tree-like tableaux to the minimal recurrent configurations of the Abelian sandpile model on these graphs. We introduce a variation on the EW-tableaux, which we call NEW-tableaux, and present bijections from these to Le-tableaux and tree-like tableaux. We also present results on various properties of and statistics on EW-tableaux and NEW-tableaux, as well as some open problems on these.

Mathematics Subject Classifications: 05A19, 05A05, 05A15, 60J10

*This work was supported by grant EP/M015874/1 from The Engineering and Physical Sciences Research Council.

†This work was supported by grant EP/M027147/1 from The Engineering and Physical Sciences Research Council.

1 Introduction and preliminaries

This paper presents a series of results on properties of the previously little studied EW-*tableaux* described by Ehrenborg and van Willigenburg [14]. Most of these results arise from a bijection we introduce between these tableaux and permutations. As a consequence of that bijection we also give a bijection between EW-tableaux and the much studied Le-tableaux, as well as to the *tree-like tableaux* of Aval, Boussicault and Nadeau [2]. We refer to the latter simply as *tree-tableaux*. Another bijection that we present between EW-tableaux and tree-tableaux leads to a bijection between tree-tableaux and the minimal recurrent configurations of the Abelian sandpile model on the Ferrers graphs described in [14], via acyclic orientations with a unique sink on these bipartite graphs (as these orientations are in one-to-one correspondence with the EW-tableaux). That bijection uses the bijection of Cori and Le Borgne [5] between the spanning trees of a graph and the recurrent configurations of the Abelian sandpile model on the graph.

We also present a variation on the EW-tableaux, which we call NEW-tableaux. In addition, we give a number of results on how properties of these two types of tableaux manifest themselves in properties of their corresponding permutations and in properties of orientations on the Ferrers graphs they are in a natural bijection with.

In [14, Cor. 4.5], Ehrenborg and van Willigenburg described a set of 2-colorings of the cells of Ferrers diagrams satisfying certain conditions (a Ferrers diagram is a two-dimensional left-aligned array of boxes, weakly decreasing in the number of boxes when going down, see Figure 2.1). These colorings have interesting properties that the authors studied in connection with other structures related to such diagrams. In particular, the set of colorings they define are in one-to-one correspondence with acyclic orientations, with a unique sink, of certain bipartite graphs, which they call Ferrers graphs. We define here a set of 0/1-fillings of Ferrers diagrams that are essentially equivalent to the definition of Ehrenborg and van Willigenburg, with one minor difference that we explain in Remark 2 below. We refer to the tableaux arising thus by the first letters of the last names of Ehrenborg and van Willigenburg.

Definition 1. An EW-*tableau* $\mathcal{T}$ is a 0/1-filling of a Ferrers diagram with the following properties:

1. The top row of $\mathcal{T}$ has a 1 in every cell.
2. Every other row has at least one cell containing a 0.
3. No four cells of $\mathcal{T}$ that form the corners of a rectangle have 0s in two diagonally opposite corners and 1s in the other two.

The *size* of an EW-tableau is one less than the sum of its number of rows and number of columns.

An example of an EW-tableau is shown in Figure 2.1.

Remark 2. Ehrenborg and van Willigenburg considered colorings corresponding to tableaux that in our definition would have a unique fixed, but arbitrary, row of all 1s, and no column with all 0s. As all our results naturally restrict to tableaux of a fixed shape, it is irrelevant where this fixed all-1s row is chosen, and as it facilitates the presentation of our results, we choose to always make this the top row. That implies, of course, that no column has all 0s. The choice of this row being arbitrary is directly linked to the fact, proved by Greene and Zaslavsky [16, Thm. 7.3], that the number of acyclic orientations, with a unique sink, of a graph is independent of which vertex is chosen as the unique sink. In fact, it is easy to show that given an acyclic orientation with a unique sink, one can reverse the direction of every edge in every directed path from an arbitrary other vertex v to the sink to obtain an acyclic orientation where v is the unique sink.

We label rows and columns on the southeast border of Ferrers diagrams underlying EW-tableaux with the numbers $0, 1, \ldots, n$, where the top row gets label 0 and the successive border edges get the remaining numbers in order, as shown in Figure 2.1.

In [14, Cor. 4.5], Ehrenborg and van Willigenburg proved the following theorem, in a different but equivalent form. That these are equivalent follows from Lemma 2.2 in [13]. The *excedance bottoms set* of an n-permutation $a_1 a_2 \ldots a_n$ is the set of indices i such that $a_i > i$. An n-permutation is a permutation of the integers $[n] = \{1, 2, \ldots, n\}$.

Theorem 3. *The* EW*-tableaux of size n with row labels $0, e_1, e_2, \ldots, e_k$ are equinumerous with n-permutations with excedance bottoms set $\{e_1, e_2, \ldots, e_k\}$.*

The proof in [14, Cor. 4.5] was, however, not bijective. One of the main results of the present paper is a bijection proving this, which leads to many new results on these tableaux, and their connections to other kinds of tableaux and other combinatorial objects. In particular, this leads to a bijection with the Le-tableaux, first introduced by Postnikov [17], and then studied in a number of papers by various authors (see, for example, [6] and [7] for references). The Le-tableaux have been shown [8, 9] to encode the PASEP, a much studied one-dimensional lattice model in statistical mechanics, and to provide a refinement of statistics on the steady state distribution of that model.

Definition 4. A Le*-tableau* $\mathcal{T}$ is a 0/1-filling of a Ferrers diagram, some of whose bottommost rows may be empty, satisfying the following properties:

1. Every column of $\mathcal{T}$ has a 1 in some cell.
2. If a cell has a 1 above it in the same column and a 1 to its left in the same row then it has a 1.

The *size* of a Le-tableau is the sum of its number of rows and number of columns.

The difference in the definitions of size for Le-tableaux compared to EW-tableaux is due to the fact that as defined both correspond to permutations whose length equals the size of the tableaux in question.

A bijection from Le-tableaux to permutations is essentially defined in [17], and explicitly described in [21], where various statistics preserved by this bijection are treated.

In particular, as was shown in [17], rows in a LE-tableau correspond to *weak excedances* in the corresponding permutation, that is, to indices i in a permutation $a_1a_2\dots a_n$ such that $a_i \geqslant i$.

A *Ferrers graph* (see [14]) is a bipartite graph whose vertices of each part are labeled $t_0, t_1, \dots, t_k$ and $b_1, b_2, \dots, b_m$, respectively, satisfying the following conditions:

1. If (t_i, b_j) is an edge and $r \leqslant i$ and $s \leqslant j$, then (t_r, b_s) is also an edge.

2. Both (t_0, b_m) and (t_k, b_1) are edges.

Given a Ferrers diagram with rows labeled from top to bottom with $t_0, t_1, \dots, t_k$ and columns labeled with $b_1, b_2, \dots, b_m$ from left to right, there is a unique Ferrers graph whose vertices are labeled with the t_i and b_j and where (t_i, b_j) is an edge if and only if the diagram has a cell in row t_i and column b_j. This correspondence is clearly one-to-one.

Lemma 5. *If an orientation of a Ferrers graph contains a cycle, then it contains a 4-cycle.*

Proof. Let G be a Ferrers graph with top vertices $t_0, t_1, \dots, t_k$ and bottom vertices $b_1, b_2, \dots, b_m$. Suppose an orientation $\mathcal{O}$ of G contains a cycle of length $2n$ for some $n > 2$. Write this cycle as $C = t_{i_1}, b_{j_1}, \dots, t_{i_n}, b_{j_n}, t_{i_1}$ where the notation means that edges are oriented $t_{i_q} \to b_{j_q}$ and $b_{j_q} \to t_{i_{q+1}}$, with the convention $i_{n+1} = i_1$. Without loss of generality, we may assume that $i_1 = \max\{0 \leqslant i \leqslant k;\, t_i \in C\}$.

Now since (t_{i_1}, b_{j_n}) is an edge of G and $i_2 \leqslant i_1$ it follows that (t_{i_2}, b_{j_n}) is also an edge of G. There are two possibilities for its orientation in $\mathcal{O}$.

1. If (t_{i_2}, b_{j_n}) is oriented $t_{i_2} \to b_{j_n}$, then $t_{i_1}, b_{j_1}, t_{i_2}, b_{j_n}, t_{i_1}$ is a 4-cycle, and the proof is complete.

2. If (t_{i_2}, b_{j_n}) is oriented $b_{j_n} \to t_{i_2}$, then $C' := t_{i_2}, b_{j_2}, \dots, t_{i_n}, b_{j_n}, t_{i_2}$ is a cycle of length $2n - 2$.

In case 2, we simply iterate the reasoning on the cycle C' until we reach a cycle of length 4, as desired. □

Since a cycle in an orientation of a Ferrers graph implies the existence of a 4-cycle, it follows that there is a one-to-one correspondence between acyclic orientations of a Ferrers graph, with a unique sink t_0, and EW-tableaux on the corresponding Ferrers diagram, as pointed out in [14]. Namely, an edge (t_i, b_j) is oriented from t_i to b_j if and only if the cell in row i and column j has a 0. Thus, conditions 1 and 2 in Definition 1 correspond to t_0 being the unique sink, and condition 3 corresponds to there being no directed 4-cycles (and therefore no directed cycles at all).

Since we have a bijection from EW-tableaux to permutations, and there exists a bijection between permutations and LE-tableaux, this induces a bijection between the two kinds of tableaux. However, the bijection we present here, to be described explicitly later, is constructed by first translating the permutation obtained from an EW-tableau, in two steps, into a different permutation. The first translation is by the following bijection,

which is a variation on the *transformation fondamentale* of Foata and Schützenberger [15]. This variation was described in the Appendix in [20] (the author was not aware of this being a variation on the TF). The *descent tops set* of a permutation $a_1a_2\dots a_n$ is the set of those letters a_i such that $a_i > a_{i+1}$ and the *descent bottoms set* consists of those a_{i+1} satisfying this. The excedance bottoms set was defined before to be the set of indices i such that $a_i > i$ and the *excedance tops set* is the set of letters a_i satisfying this.

Proposition 6. *Let $\pi = a_1a_2\dots a_n$ be an n-permutation. The following defines a bijection*[1] $\mathrm{DE} : \mathcal{S}_n \to \mathcal{S}_n$ *such that* $\mathrm{DE}(\pi) = b_1b_2\dots b_n$*, where each b_i is determined as follows: Let $a_0 = 0$.*

1. *If $a_i > a_j$ for some $j > i$ then $b_{a_{i+1}} = a_i$, that is, a_i is then in place a_{i+1} in* $\mathrm{DE}(\pi)$.
2. *If $a_i < a_j$ for all $j > i$, find the rightmost letter smaller than a_i. If this letter is a_k then $b_{a_{k+1}} = a_i$, that is, then a_i is in place a_{k+1} in* $\mathrm{DE}(\pi)$.

The excedance bottoms set of $\mathrm{DE}(\pi)$ *equals the descent bottoms set of π.*

As an example, let $\pi = 361542$. Then $\mathrm{DE}(\pi) = 641523$: Since 3 has a smaller letter to its right 3 goes to the place determined by the letter following 3, namely 6. Since 1 has no smaller letter to its right we find the rightmost letter smaller than it, which is defined to be $a_0 = 0$, so 1 goes to place $a_1 = 3$, etc.

The second transformation, which we apply to the image of DE, is the cyclic right shift $\mathcal{R}$, that is, the bijection that maps $a_1a_2\dots a_n$ to $a_na_1a_2\dots a_{n-1}$. The reason for this is that we can associate various properties of an EW-tableau $\mathcal{T}$ to properties of $\sigma = \mathrm{DE}(\Psi(\mathcal{T}))$, where Ψ is the bijection taking EW-tableaux to permutations, to be defined in the next section. By right shifting σ we turn excedances into weak excedances, in addition to 1 always becoming a weak excedance bottom when we move a_n, never an excedance top, to the first place, so we get exactly one more weak excedance in $\mathcal{R}(\pi)$ than we have excedances in π.

The paper is organized as follows: In Section 2 we describe our bijection from EW-tableaux to permutations. In Section 3 we define NEW-tableaux and describe some properties of the permutations obtained from those tableaux when applying to them the bijection from EW-tableaux to permutations. We describe a bijection between EW-tableaux and Le-tableaux in Section 4, a bijection between NEW-tableaux and Le-tableaux in Section 5 and a bijection between EW-tableaux and tree-tableaux in Section 6, the last of which can be easily modified to become a bijection between NEW-tableaux and tree-tableaux. In Section 7 we present a connection between tree-tableaux and certain spanning trees of the underlying Ferrers graph, and a connection between EW-tableaux and minimal recurrent configurations of the Abelian sandpile model, and use this to give a new bijection between tree-tableaux and EW-tableaux, which is different from the bijection in Section 6. Finally, in Section 8 we give several results on properties and statistics of EW- and NEW-tableaux and their corresponding permutations, and a few open problems on these.

[1]The name DE refers to the map taking descent tops/bottoms to excedance tops/bottoms.

In a forthcoming paper [12], we describe a certain "decoration" of the EW-tableaux and their corresponding permutations. There is a one-to-one correspondence between these decorated objects and all recurrent configurations of the Abelian sandpile model on Ferrers graphs. This complements and extends the work of Dukes and Le Borgne [11] (see also [3]) on the sandpile model on the complete bipartite graph, which is the special case of Ferrers graphs corresponding to rectangular Ferrers diagrams.

2 The bijection from EW-tableaux to permutations

We now define a map Ψ from EW-tableaux of size n to n-permutations, which we then show to be well defined, and subsequently that it is a bijection. See Figure 2.1 and Example 8 for an example of this.

Definition 7. Given an EW-tableau $\mathcal{T}$, delete all entries from its top row and disregard the label 0 in what follows. We construct a permutation $\pi = \Psi(\mathcal{T})$ by repeating the following two steps until all entries of $\mathcal{T}$ have been deleted:

1. Write down labels of columns with no 1s, one after the other from right to left, and delete the entries from each of these columns.
2. Write down labels of rows with no 0s, one after the other from bottom to top, and delete the entries from each of these rows.

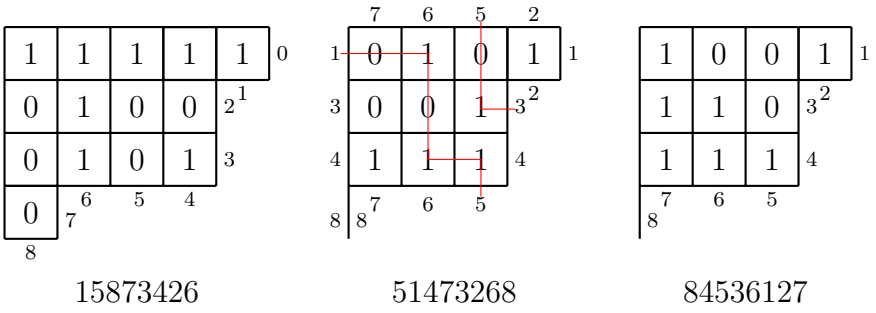

Figure 2.1: Example of an EW-tableau (left), a LE-tableau (center) and a NEW-tableau, with their corresponding permutations. Shown (in red) are two of the paths in the LE-tableau determining its corresponding permutation (5 in the first place, 3 in the fifth, etc.), see [21]. These paths are determined as follows: From a label ℓ on the top or left border, head into the tableau. When you hit a 1, turn right if you are heading down, down if you are heading right. The label you hit when exiting the tableau is the letter in place ℓ in the permutation.

Example 8. Figure 2.1 shows an example of the bijection Ψ in Definition 7 for the EW-tableau (on the left). Delete the entries in the top row. We first read the labels of columns with no 1s remaining, in increasing order: $1, 5, 8$; and delete the entries in these

columns. Then the labels of rows with no remaining 0s, in decreasing order, namely 7 and 3, deleting the entries in these rows, and then repeat. The same algorithm is applied to the NEW-tableau (on the right), see Section 3, except that we don't delete the top row. In this case, however, there will be no columns without a 1, so we first read the rows 8 and 4, then column 5 and so on.

Lemma 9. *At each stage of the process in Definition 7 there is an unread column with no 1s or an unread row with no 0s.*

Proof. If all columns have been read then no row has a 0, and if all rows have been read then no column has a 1. When there are both rows and columns unread, if each of those rows has a 0 and each of those columns has a 1, then we can go from a 0 in a row to a 1 in that 0's column, then to a 0 in that 1's row, ad infinitum, which would imply the corresponding Ferrers graph had a directed cycle, contradicting Lemma 5 and its subsequent paragraph. □

Lemma 9 shows that performing the process described in Definition 7 we will read all the labels of rows and columns of $\mathcal{T}$, and thus eventually produce an n-permutation. The map Ψ is therefore well defined. We now show that it is a bijection, but we first need a lemma.

Lemma 10. *When constructing the permutation $\pi = \Psi(\mathcal{T})$ of an* EW*-tableau $\mathcal{T}$ by the algorithm described in Definition 7, we first record a nonempty sequence of column labels in increasing order, then a nonempty sequence of row labels in decreasing order and so on, alternatingly. Moreover, a row label is always preceded by a larger letter in π and a column label is never preceded by a larger letter.*

Proof. After having deleted all entries from the top row of an EW-tableau $\mathcal{T}$, since no other row can have all 1s and thus be 0-free, there must be a 1-free column, by Lemma 9. On the first pass of part 1 in Definition 7, we read all the 1-free columns, in order of increasing labels and delete their entries. Having done that, there must now be a 0-free row, and we read the labels of all the 0-free rows, in decreasing order, according to 2 in Definition 7, so that if a row label ℓ is read immediately after another row label m we must have $m > \ell$.

If a row label ℓ is immediately preceded by a column label c, then c was the last column label read on that pass of column labels, and therefore the largest of those labels. The label ℓ read just after that must belong to a row that had a 0 removed on the last pass of column labels read. Therefore, ℓ cannot be larger than c, for then it would be larger than all the column labels read on the preceding pass, implying that the row in question wasn't long enough to have had a column with a label of c or smaller. Thus, a row label is always preceded by a larger letter in $\pi = \Psi(\mathcal{T})$.

An argument analogous to the above shows that, except for the first column label read, a column label is always preceded by a smaller letter in $\pi = \Psi(\mathcal{T})$, proving the claim. □

The following is an immediate consequence of Lemma 10.

Proposition 11. *The labels of rows of an* EW*-tableau* $\mathcal{T}$ *are exactly all the descent bottoms in* $\pi = \Psi(\mathcal{T})$*, and thus exactly all the excedance bottoms of* $\mathrm{DE}(\pi)$*.*

Proposition 12. *The map* Ψ *is a bijection.*

Proof. We show first that Ψ is an injection. If two tableaux $\mathcal{T}$ and $\mathcal{S}$ yield the same permutation then they must have the same shape, by Proposition 11, and the same set of 1-free columns (after clearing the top row). Once the entries of these columns have been deleted in both, $\mathcal{T}$ and $\mathcal{S}$ must have the same 0-free rows, since the labels of the same rows will be read on the next pass of the algorithm in Definition 7. But that implies that these rows were identical in $\mathcal{T}$ and $\mathcal{S}$, since both had 0s removed from all the same columns on the first pass. By induction this must then apply to all columns and rows.

By Lemma 3, applied to all possible sets of row labels, the set of all EW-tableaux of size n is equinumerous with n-permutations, and thus injectivity of Ψ implies that it is also surjective. □

It is straightforward, from the algorithm in Definition 7, to verify the following.

Proposition 13. *The following describes how to construct the inverse image* $\mathcal{T}$ *under* Ψ *of a permutation* $\pi = a_1 a_2 \dots a_n$*:*

The set of descent bottoms of π*, together with* n*, determines the shape of* $\mathcal{T}$ *and the labels of rows and columns. Fill the top row of* $\mathcal{T}$ *with 1s. Now, since the letter* a_1 *is the first letter of the permutation and thus not a descent bottom, it is necessarily the label of a column; fill the empty cells of that column with 0s. If* a_2 *labels a column, fill the empty cells of that column with 0s, otherwise fill the empty cells of the row labeled by* a_2 *with 1s. Now repeat this for each* a_i*.*

This, in turn, obviously leads to the following description of the inverse of Ψ:

Corollary 14. *Let* $\pi = a_1 a_2 \dots a_n = \Psi(\mathcal{T})$*. If* r *is the label of a row and* c *the label of a column in* $\mathcal{T}$ *then the cell in row* r *and column* c *has a 1 if and only if* r *precedes* c *in* π*.*

An *up-down run* in an n-permutation π is a maximal sequence of consecutive letters in π that is either increasing or decreasing, when we prepend 0 to π so that the first run is always increasing. For example, $\pi = 51368427$ has up-down runs 05, 51, 1368, 842, 27. The *run-decomposition* of a permutation is its sequence of up-down runs where the first letter of each run has been removed. Thus, π has run-decomposition 5-1-368-42-7. The number of up-down runs in π is thus one greater than the sum of the numbers of peaks and valleys in 0π. The distribution of the numbers of up-down runs in permutations is found in sequence A186370 in [1]. The following lemma is now straightforward to prove, the key observation being from Lemma 10 that the first row label read after reading a sequence of column labels must be smaller than the last column label read, and that the first column label read after reading row labels must be greater than the last row label read.

Lemma 15. *The blocks in the run-decomposition* R *of a permutation* $\pi = \Psi(\mathcal{T})$*, for an* EW*-tableau* $\mathcal{T}$*, correspond to the alternating sequences of column and row labels as they are read in constructing* π *from* $\mathcal{T}$*.*

3 NEW-tableaux

If we reflect an EW-tableau along the NW-SE diagonal and turn 0s into 1s and vice versa, we get a tableau whose leftmost column is all 0. If we remove that leftmost column, but keep all its rows, the bottom few of which may now be empty, we get a tableau satisfying the conditions in the following definition. Note that, by definition, every non-top row in an EW-tableau has a 0, which is reflected in condition 1 in the following definition.

Definition 16. A NEW-*tableau*[2] $\mathcal{T}$ is a 0/1-filling of a Ferrers diagram, possibly with some empty bottom rows, satisfying the following conditions:

1. Every column of $\mathcal{T}$ has a 1.
2. No four cells of $\mathcal{T}$ that form the corners of a rectangle have 0s in two diagonally opposite corners and 1s in the other two.

The *size* of a NEW-tableau is the sum of its number of rows and number of columns.

It is easy to see that Lemma 9 also applies to NEW-tableaux, and therefore we can apply the map Ψ to such a tableau $\mathcal{T}$, yielding a permutation whose length equals the size of $\mathcal{T}$. The proof of Proposition 12 also goes through when Ψ is applied to a NEW-tableau, with the only change that we observe, for the last paragraph of the proof, that NEW-tableaux of size n are equinumerous with EW-tableaux of size n, by the discussion preceding Definition 16, and therefore with n-permutations.

The properties of a permutation $\pi = \Psi(\mathcal{T})$ for a NEW-tableau $\mathcal{T}$ are slightly different from the case of EW-tableaux. Note that Lemma 9 ensures that some row in any NEW-tableau $\mathcal{T}$ is all 1s, since every column of $\mathcal{T}$ has a 1.

Theorem 17. *Let $\pi = \Psi(\mathcal{T})$ for a* NEW-*tableau $\mathcal{T}$. Then π and $\mathcal{T}$ have the following properties:*

1. *The label of the lowest (possibly empty) 0-free row of $\mathcal{T}$ equals the first letter in π, and thus is the place of the letter 1 in* $\mathrm{DE}(\pi)$.
2. *The labels of other rows in $\mathcal{T}$ are descent bottoms in π, and thus excedance bottoms of* $\mathrm{DE}(\pi)$.
3. *If r is the label of a row and c the label of a column in $\mathcal{T}$ then the cell in row r and column c has a 1 if and only if r precedes c in π.*

Proof. Since every column in $\mathcal{T}$ has a 1, we will inevitably first read the label of a row when performing the algorithm described in Definition 7. Since row labels are read from bottom to top, we first read the label of the lowest row with no 0s (and the first letter in a permutation π is always the place of 1 in $\mathrm{DE}(\pi)$). Apart from this, the argument in the proof of Lemma 10 applies here.

[2]NEW stands for New EW-tableau.

Since the bijection Ψ is defined in the same way for NEW-tableaux as for EW-tableaux, Proposition 13 and Corollary 14 apply here, except that we don't start by filling the top row with 1s, and the first letter of π is always the label of a row. □

4 A bijection between EW- and Le-tableaux

Both EW- and Le-tableaux of size n map bijectively to n-permutations. Moreover, the shape of an EW-tableau $\mathcal{E}$ is determined by the excedance bottoms set of the permutation $\mathrm{DE}(\Psi(\mathcal{E}))$, and the shape of a Le-tableau $\mathcal{T}$ is determined by the weak excedance bottoms of the permutation $\Phi(\mathcal{T})$ described in [21] (see the caption of Figure 2.1). As pointed out before, by cyclically right shifting a permutation $\pi = a_1 a_2 \dots a_n$ to obtain a permutation $\pi' = \mathcal{R}(\pi)$, the set of excedance tops in π becomes the set of weak excedance tops in π', except that a_n, which is never an excedance top goes to the first place in π', adding one weak excedance top in π'. Thus, composing these bijections we get a bijection Δ from the set of EW-tableaux of size n to the set of Le-tableaux of size n, namely,

$$\Delta = \Phi^{-1} \circ \mathcal{R} \circ \mathrm{DE} \circ \Psi\,. \tag{4.1}$$

Since all of the bijections in this composition have well understood inverses (see [19] for the inverse of DE), it is clear how to define Δ^{-1}. We now describe how to effect the bijection Δ^{-1} directly from Le-tableaux to EW-tableaux, and later give a direct description of Δ.

Let $\mathcal{L}$ be a Le-tableau. We construct the corresponding EW-tableau $\mathcal{E} = \Delta^{-1}(\mathcal{L})$. The shape of $\mathcal{E}$ is obtained by taking the shape of $\mathcal{L}$ and adding a column to its left border, of the same length as that border, which may have empty rows at the bottom. We label the rows and columns of $\mathcal{L}$ and $\mathcal{E}$ in the following ways:

Starting at the top right corner of $\mathcal{L}$, label the border edges increasingly, beginning with 1, as you follow the path down to the bottom left. We call these labels the *r-labels* of $\mathcal{L}$. Repeat the same process for $\mathcal{E}$ but begin with 0. Label the left and top edges of $\mathcal{L}$ by $e-1$, where e is the r-label at the other end of that row or column, replace 0 by n, and call these the *ℓ-labels* of $\mathcal{L}$. See Figure 4.1 for an example of this labeling.

The paths referred to in the description below are essentially the same paths as in [21, Section 2], except that our paths here go in the "backwards" direction compared to [21]. Also, instead of having our paths go east or south on their last leg they end by going west or north, that is, our paths terminate at the opposite end of the row or column compared to the paths in [21], as shown in Figure 4.1.

We now present the method for filling $\mathcal{E}$ based on the filling of $\mathcal{L}$. This also returns the permutation $\pi = \Psi(\mathcal{E})$.

Algorithm 18.

1. Set all r-labels of $\mathcal{L}$ as unvisited, set π as the empty word and fill the top row of $\mathcal{E}$ with 1s.

2. Let i be the smallest unvisited r-label of $\mathcal{L}$. If there are no unvisited r-labels left then stop.

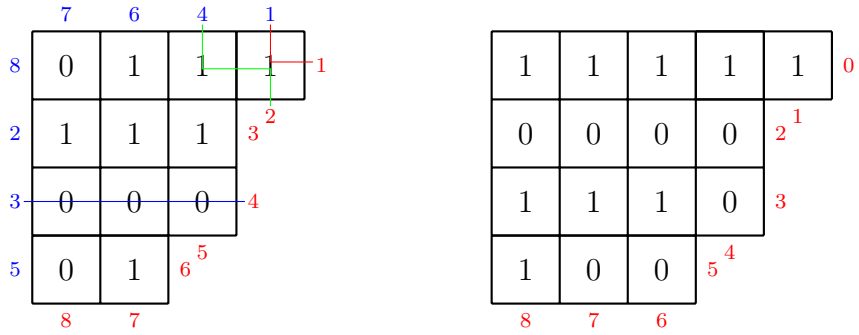

Figure 4.1: An example of a LE-tableau (left) and corresponding EW-tableau (right) along with their r-labels (red) and ℓ-labels (blue), used in Algorithm 18. Also shown are their corresponding permutations, and the intermediate permutation DE(14367582), which is a cyclic left shift of 51842736. In the LE-tableau, the first three paths considered by Algorithm 18, starting at r-labels 1, 2 and 4, respectively, are shown in red, green and blue, respectively.

3. Mark i as visited and follow the path from the label i until you hit an ℓ-label, call it j. This is the path where when you hit a 1 you turn west if you were going north and north if you were going west.

4. If j is a column r-label of $\mathcal{E}$ fill any empty cells of $\mathcal{E}$ in that column with 0s. If j is a row r-label of $\mathcal{E}$ fill the empty cells in that row with 1s.

5. Add j to the end of π. If j is an unvisited r-label of $\mathcal{L}$, then set $i = j$ and go to Step 3, otherwise go to Step 2.

A tedious but straightforward argument, using the above algorithm for constructing $\mathcal{E} = \Delta^{-1}(\mathcal{L})$, leads to the following description of how a LE-tableau $\mathcal{L} = \Delta(\mathcal{E})$ is constructed directly from the permutation $\pi = \Psi(\mathcal{E})$. That in turn leads to the following direct construction of $\mathcal{L}$ from $\mathcal{E}$, since π simply records the reading of columns and rows of $\mathcal{E}$.

Let $m(i)$ be the position of the letter i in $\mathrm{DE}(\Psi(\mathcal{E}))$. Equivalently, $m(i)$ can be read directly from $\pi = \Psi(\mathcal{E})$, after we prepend 0 to π, as follows: Recall that the right-to-left minima of π are the letters with no smaller letters to their right. If the letter $i \geqslant 1$ in π is a right-to-left minimum let $m(i)$ be the letter immediately following the next right-to-left minimum to the left of i in π, otherwise let $m(i)$ be the letter immediately to the right of i in π.

The shape of $\mathcal{L}$ is obtained by deleting the leftmost column of $\mathcal{E}$, but keeping the possibly empty bottom rows resulting from that. We fill $\mathcal{L}$ as follows, with i initialized to 1, n the size of $\mathcal{L}$, and if $r = 1$ we let $r - 1 = n$:

Algorithm 19. 1. Enter $\mathcal{L}$ from the r-label i, let $d \in \{N, W\}$ be the direction of travel and (r, c) the current cell, where r and c are the row and column r-labels of that cell.

2. If (r, c) is filled, then go to Step 4. If (r, c) is empty, then fill it with $f(r, c)$ as follows:

   - When $m(i) = r - 1$, if $d = W$ then $f(r, c) = 0$, else $f(r, c) = 1$.
   - When $m(i) = c - 1$, if $d = N$ then $f(r, c) = 0$, else $f(r, c) = 1$.
   - When $d = W$, $m(i) \neq r - 1$ and there is a 1 in the column above (r, c), then $f(r, c) = 1$.
   - When $d = W$, $m(i) \neq c - 1$ and all cells in the column above (r, c) are 0s, then $f(r, c) = 0$.

3. If (r, c) contains a 0 move to the next cell in direction d and if (r, c) contains a 1 move to the next cell in the "opposite" direction to d, that is, north if $d = W$ and west if $d = N$.

4. If we haven't reached the edge of $\mathcal{L}$ go to Step 2. If we reach the edge of $\mathcal{L}$, then increase i by 1, and if $i \leqslant n$ go to Step 1, otherwise stop.

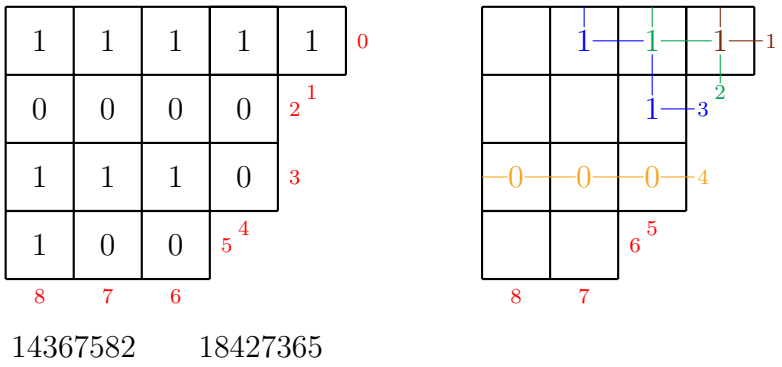

Figure 4.2: The EW-tableau $\mathcal{T}$, and its associated permutations $\Psi(\mathcal{T})$ (left) and $\mathrm{DE}(\Psi(\mathcal{T}))$ (right), from Figure 4.1, and the corresponding LE-tableau partially filled by applying Algorithm 19.

5 A bijection between NEW- and LE-tableaux

It is, of course, possible to define a bijection between NEW-tableaux and LE-tableaux via permutations, in the exact same way as we did in the case of EW-tableaux and LE-tableaux. However, a variation on that turns out to have nicer properties, namely to preserve shape exactly. As noted in Theorem 17, the row labels of a NEW-tableau $\mathcal{N}$ consist of the (number of the) place of 1 in the permutation $\pi = \mathrm{DE}(\tau)$ and the excedance bottoms of π, where τ is the permutation read from $\mathcal{N}$ according to Definition 7. We call this set of excedance bottoms together with the place of 1 the *augmented excedance bottoms set* of π.

If n is the length of a permutation π then the *cyclic down-shift* $\mathcal{S}$ subtracting 1 from each entry of π and changing the resulting 0 to n takes the augmented excedance bottoms set of π to the set of weak excedance bottoms of π. As an example, $3164275 \mapsto 2753164$, where $\{1, 2, 3, 6\}$ is the augmented excedance bottoms set of the first permutation and therefore the set of weak excedance bottoms of the second.

Thus, according to Theorem 17, and the fact that the row labels of an EW-tableau $\mathcal{T}$ are the excedance bottoms of the corresponding permutation $\mathrm{DE}(\Psi(\mathcal{T}))$, we can define a bijection from NEW-tableaux to LE-tableaux as the following composition (cf. Equation 4.1):

$$\Gamma = \Phi^{-1} \circ \mathcal{S} \circ \mathrm{DE} \circ \Psi\,. \tag{5.1}$$

With one modification the algorithm in 18 goes through to give a direct description of the above bijection. Namely, instead of shifting the labels on the top and left border of a LE-tableau $\mathcal{L}$, we apply the transformation $\mathcal{S}$ described above to the labels on the right of rows and bottom of columns. We show an example of this map in Figure 5.1.

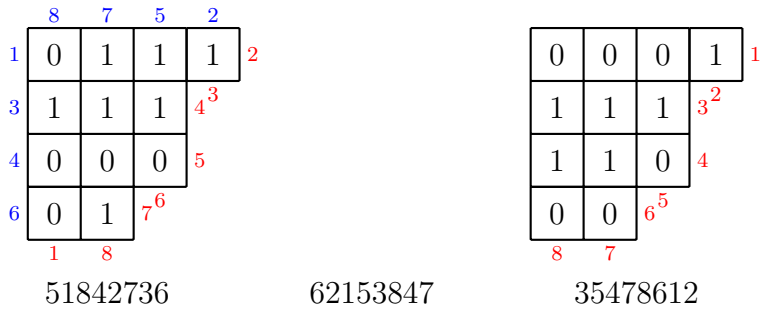

Figure 5.1: An example of a LE-tableau (left) and corresponding NEW-tableau (right) along with their r-labels (red) and ℓ-labels (blue). Also shown are their corresponding permutations, and the intermediate permutation $\mathrm{DE}(35478612)$, which is a cyclic up-shift of 51842736.

6 A bijection between (N)EW- and tree-tableaux

In [2], Aval, Boussicault and Nadeau introduced what they call *tree-like tableaux*, Ferrers diagrams where some cells have a dot and others are empty. We will here refer to these tableaux simply as *tree-tableaux*.

Definition 20. A *tree-tableau* is a filling of a Ferrers diagram with the following properties:

1. The top left cell has a dot.
2. Every other cell has a dot above it (in the same column) or left of it (in the same row), but not both.
3. Every column and every row has a dot.

The *size* of a tree-tableau is one less than the sum of its number of rows and number of columns.

Tree-tableaux have some very nice properties and strong connections to LE-tableaux and the *alternative tableaux* defined by Viennot [22]. A bijection between tree-tableaux and LE-tableaux can be inferred from [2, Prop. 3]. We describe this bijection below, and how to translate it into a bijection to EW- and NEW-tableaux (see Figure 6.1) but first we introduce some notation.

Define a dot in a tree-tableau as *left-free* if it has no dots to its left in its row and *up-free* if it has no dots above it in its column. By the definition of a tree-tableau every dot is either up-free or left-free, except the top left cell, which is both. Moreover, every row contains exactly one left-free dot and every column contains exactly one up-free dot. A cell has a dot *weakly above* it if it has a dot or there is a dot above it in the same column. Similarly a cell has a dot *weakly right* of it if it has a dot or there is a dot to its right in the same row. A 0 in a LE-tableau is *restricted* if it has a 1 above it in the same column (see [9]).

The map T from a tree-tableau $\mathcal{T}$ to a LE-tableau can be deduced from [2, Prop. 3] and described as follows: Change each left-free dot to a 0 and fill all cells to its left with 0s. Change every up-free dot to a 1 and fill all cells above it with 0s. Finally, delete the leftmost column of $\mathcal{T}$ and fill all remaining empty cells with 1s.

The inverse map Λ, from a LE-tableau $\mathcal{L}$ to a tree-tableau, can be described as follows: Add a column to the left of $\mathcal{L}$ with a 1 in the top cell and all 0s below. Turn every 1 that is highest in its column into a dot, turn the rightmost restricted 0 in each row to a dot and then leave all remaining cells empty.

We now define a map F from tree-tableau to LE-tableau, which we show to be just a simpler description of T. Let $\mathcal{T}$ be a tree-tableau and make the following changes:

- If a cell has a dot to its left and a dot weakly above, fill it with a 1.
- Otherwise fill the cell with a 0.
- Finally, delete the leftmost column to obtain $\mathrm{F}(\mathcal{T})$.

Lemma 21. *Given any tree-tableau $\mathcal{T}$ we have $\mathrm{T}(\mathcal{T}) = \mathrm{F}(\mathcal{T})$.*

Proof. Consider a cell c in $\mathcal{T}$ and let $\mathrm{T}(c)$ and $\mathrm{F}(c)$ be the corresponding cells in $\mathrm{T}(\mathcal{T})$ and $\mathrm{F}(\mathcal{T})$, respectively. We show that $\mathrm{T}(c)$ and $\mathrm{F}(c)$ are filled with the same value for every cell in $\mathcal{T}$.

Firstly suppose c has no dot to the left, so $\mathrm{F}(c)$ contains a 0. There is a left-free dot in cell a that is weakly right of c, as every row has exactly one left-free dot. Therefore, $\mathrm{T}(a)$ and all cells to the left are filled with 0s, so $\mathrm{T}(c)$ contains 0. Next suppose c has a dot to the left but no dot weakly above, so $\mathrm{F}(c)$ contains a 0. So there is an up-free dot (strictly) below c in $\mathcal{T}$, which implies $\mathrm{T}(c)$ contains a 0.

Finally suppose c has a dot to the left and a dot weakly above, so $\mathrm{F}(c)$ contains a 1. If c contains a dot, then the dot must be up-free because it is not left-free as there is a

cell to the left with a dot, which implies $\mathrm{T}(c)$ contains a 1. If c does not contain a dot then the left-free dot of that row must be left of c, and the up-free dot of that column must be above c, which implies $\mathrm{T}(c)$ is one of the remaining empty cells that get filled with a 1. □

In Section 4 we defined a bijection Δ^{-1} from LE-tableaux- to EW-tableaux. Composing this bijection with the bijection between tree-tableaux and LE-tableaux gives a bijection from tree-tableaux to EW-tableaux. The map from a LE-tableau $\mathcal{L}$ to an EW-tableau $\mathcal{E}$ is given by Algorithm 18 by tracing paths through $\mathcal{L}$, from an r-label to an ℓ-label. We can define a path through a tree-tableau which changes direction in a cell c if $\mathrm{F}(c)$ contains a 1. So the paths through the tree-tableau are identical to the paths through the corresponding LE-tableau, apart from an additional horizontal step through the cell in the extra leftmost column, regardless of its content. A path through a tree-tableau thus changes direction in a cell c according to the following rule:

Change direction if and only if c has a dot to the left and a dot weakly above. (6.1)

The path through a tree-tableau beginning at r-label i thus ends at the same ℓ-label as the path through the corresponding LE-tableau. We can therefore apply Algorithm 18 to trace these paths through a tree-tableau to get a map from tree-tableaux to EW-tableaux, see Figure 6.1. By the last paragraph of Section 5, we can thus also use Algorithm 18 to produce a bijection from tree-tableaux to NEW-tableaux.

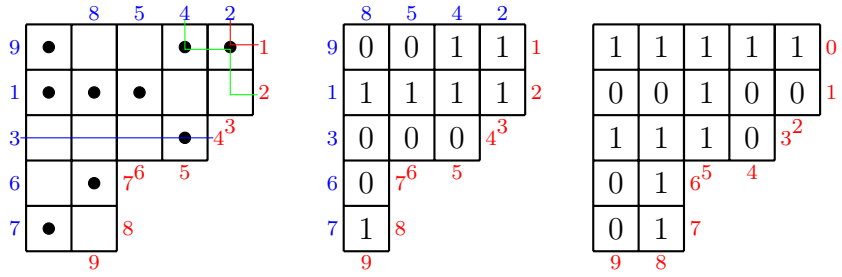

Figure 6.1: An example of a tree-tableau (left) and corresponding LE-tableau (center) and EW-tableau (right). In the tree-tableau, the first three paths given by rule (6.1), starting at r-labels 1, 2 and 4, are shown in red, green and blue, respectively.

7 Tableaux, spanning trees and the Abelian sandpile model

In this section we exhibit a connection between the tree-tableaux and certain spanning trees of the underlying Ferrers graph. We combine this with the work in [5] to obtain a new bijection between tree-tableaux and EW-tableaux. This bijection is different from that of Section 6.

7.1 EW-tableaux and minimal recurrent configurations of the Abelian sandpile model

The Abelian sandpile model (ASM) is a dynamic process on a graph $G = (V \cup \{s\}, E)$ with a special distinguished vertex s called the sink. More precisely, it is a Markov chain on the set of *configurations* on G. We consider the ASM on Ferrers graphs. Let $G(F)$ be a Ferrers graph, with vertex set $0, 1, \ldots, n$ where we identify a vertex with the label of its row or column in the corresponding Ferrers diagram. We will always take vertex 0, the label of the top row of the diagram, to be the sink of the ASM.

A *configuration on* $G(F)$ is a vector $c = (c_1, \ldots, c_n)$ of nonnegative integers. We think of c_i as the number of "grains of sand" at the vertex i. Given a configuration c, we say that the vertex i is *stable* if $c_i < d_i$, where d_i is the degree of vertex i, that is, the number of cells in the row or column labeled i in the Ferrers diagram. A configuration is stable if all its vertices are stable.

Unstable vertices *topple*, sending one grain to each of their neighbors. This may cause some of these neighbors to become unstable themselves, and topple in turn. The sink never topples, and can be viewed as absorbing grains. It is possible to show (see for instance [10, Section 5.2]) that starting from any configuration c and toppling unstable vertices, one eventually reaches a stable configuration c'. Moreover, c' does not depend on the order in which vertices are toppled in this sequence. We call c' the *stabilization* of c and write $c' = \sigma(c)$.

For a configuration c, define $\tilde{c}$ by

$$\tilde{c}_i := \begin{cases} c_i + 1, & \text{if } i \text{ is a column label,} \\ c_i, & \text{otherwise.} \end{cases}$$

A configuration c is *recurrent* if $\sigma(\tilde{c}) = c$. Moreover, in the stabilization of $\tilde{c}$, every vertex i for $1 \leqslant i \leqslant n$ topples exactly once [10, Section 6.1]. A recurrent configuration c is *minimal* if it is minimal with respect to the total number of grains $\sum_{i=1}^{n} c_i$. The set of minimal recurrent configurations of a graph G is denoted $\mathrm{Rec}^m(G)$.

Let F be a Ferrers diagram. To an EW-tableau $\mathcal{E}$ of shape F, we associate a configuration $c = c(\mathcal{E})$ on the Ferrers graph $G = G(F)$ as follows:

- If i is the label of a column C, then c_i is the number of 0s in C.
- If i is the label of a row R, then c_i is the number of 1s in R.

Since each row of an EW-tableau contains at least one 0, and each column at least one 1 (in the top row), c is stable by construction. Combining the results from [18, Cor. 3] and Section 1 on acyclic orientations, we get the following result, where $\mathrm{EWT}(F)$ is the set of EW-tableaux of shape F.

Theorem 22. *The map* $\mathrm{EWT}(F) \longrightarrow \mathrm{Rec}^m(G)$ *that maps* $\mathcal{E}$ *to* $c = c(\mathcal{E})$ *is a bijection. Moreover, let* $\tau = \Phi(\mathcal{E})$ *be the permutation obtained by Definition 7. Then, starting from the configuration* $\tilde{c}$*, the vertices of* G *can be toppled in the order of* τ*, read from left to right, to reach the configuration* c*.*

Proof. Let $G = G(F)$ be a Ferrers graph. Write $O(G)$ for the set of acyclic orientations of G where the vertex labeled 0 corresponding to the top row of F is the unique sink. For an orientation $\mathcal{O} \in O(G)$, we define a configuration $c = c(\mathcal{O})$ on G by $c_i = \text{in}_i$, the number of incoming edges to the vertex i in $\mathcal{O}$. In terms of the ASM, and after switching the orientations of all edges, Corollary 3 in [18] can be restated as follows: The map $O(G) \longrightarrow \text{Rec}^m(G)$ that maps $\mathcal{O}$ to $c = c(\mathcal{O})$ is a bijection. Combined with the observation of a one-to-one correspondence between $\text{EWT}(F)$ and $O(G(F))$ in Section 1, this yields the desired bijection.

The toppling order of $\tilde{c}$ is the same as the *firing sequence* in [18]. For an orientation $\mathcal{O}$, this is given by removing the sink and all incoming edges, recording all vertices that are now sinks, removing these, and iterating. This is the exact analogue of the construction of the bijection Ψ in Definition 7, though in the toppling order the order in which we record vertices in each block in the run-decomposition of τ is unimportant. □

7.2 Tree-tableaux and spanning trees

Let F be a Ferrers diagram, and $G = G(F)$ the corresponding Ferrers graph. A *dotted tableau* is a filling of F where each cell is either empty or contains a single dot. A tree-tableau is thus an example of a dotted tableau. To any dotted tableau $\mathcal{T}$ we may associate a subgraph $S = S(\mathcal{T})$ of G by keeping the edges corresponding to dotted cells, and removing those corresponding to empty cells. This clearly gives a bijective correspondence between dotted tableaux and subgraphs of G. Now let $<_E$ be a total order on the edges of G, and S be a spanning tree of G. An edge $e \notin S$ is *externally active* if it is the minimal edge in the unique cycle of $S \cup \{e\}$ for the order $<_E$. The external activity $e(S)$ of a spanning tree S is its number of externally active edges. In this subsection we identify a cell in a Ferrers diagram with its corresponding edge in the Ferrers graph.

Definition 23. Let $G = G(F)$ be a Ferrers graph, and $<_E$ a total order on its edges or, equivalently, on the cells of F. We say that $<_E$ is *F-compatible* if it is increasing in every row from left to right, and every column from top to bottom, in F.

Theorem 24. *Let $G = G(F)$ be a Ferrers graph, and $<_E$ an F-compatible order. Let $\mathcal{T}$ be a dotted tableau of shape F. Then $\mathcal{T}$ is a tree-tableau if and only if $S(\mathcal{T})$ is a spanning tree with external activity 0.*

Proof. First suppose that $\mathcal{T}$ is a tree-tableau, and let $S = S(\mathcal{T})$. We first show that S contains no cycles. Assume that S does contain a cycle C, and consider the lowest row in $\mathcal{T}$ that contains a dotted cell belonging to C, and take c to be the rightmost dotted cell in that row belonging to C. Then each vertex of c must be connected by an edge to at least one other vertex in the cycle C, and by construction this means that c must have a dotted cell above it in its column, and to its left in its row, which contradicts the definition of a tree-tableau. Thus S contains no cycles. Moreover, from [2], we have that the number of dotted cells in $\mathcal{T}$, that is, of edges in S, is one less than the number of vertices of G, and thus S is a spanning tree.

We now need to show that $e(S) = 0$. First we will prove two useful lemmas. Since S is a spanning tree, for any pair of vertices v, w in the graph G, there is a unique path from v to w, which we denote $v \to w$.

Lemma 25. *Let v be a vertex of S. Then on the unique path $v \to 0$ in S, the row labels of vertices appear in decreasing order, and the column labels in increasing order.*

Proof. Given a tree-tableau $\mathcal{T}$ and a vertex v of $S(\mathcal{T})$ we obtain the path from $v \to 0$ via the following steps,

- if v is a row of the Ferrers diagram, set v' to be the leftmost column of $\mathcal{T}$ that contains a dot in the row v;
- if v is a column of the Ferrers diagram, set v' to be the topmost row of $\mathcal{T}$ that contains a dot in the column v;

and iterating the above on v' until reaching the top row, which is the vertex 0. Since going from bottom to top decreases row labels, and going from right to left increases column labels, we have the desired result. □

Lemma 26. *Let (i, j) and (i', j') be two edges of G such that $i' \leqslant i$ and $j' \geqslant j$, that is, the cell (i', j') is above and to the left of the cell (i, j). Then $(i', j') \leqslant_E (i, j)$.*

Proof. Since F is a Ferrers diagram, the edge (i', j) is in G, and the corresponding cell is to the right of (i', j') in the same row and above (i, j) in the same column. The order $<_E$ being F-compatible, this implies $(i', j') \leqslant_E (i', j) \leqslant_E (i, j)$ as desired. □

We return to the proof that $e(S) = 0$. Let $e = (i, j)$ be an edge that isn't in S, where i, resp. j, is the label of the row, resp. column, of the cell e. The unique cycle C in $S \cup \{e\}$ is obtained by taking the union of the paths $i \to 0$ and $j \to 0$, removing edges that appear in both paths, and combining with the edge e. Let i' be the first row label encountered on the path $j \to 0$ such that $i' \leqslant i$, and j' the column label of the vertex immediately preceding i' on this path. By Lemma 25, $j' \geqslant j$. If (i', j') isn't on the path $i \to 0$, then it belongs to C, and by Lemma 26, we have $(i', j') <_E (i, j)$. The inequality is strict since (i, j) is not on the path $j \to 0$ (it isn't in S). Thus (i, j) is not externally active, since C contains an edge that is strictly smaller in $<_E$. If the path $i \to 0$ does go through j' and i' in that order (it must go through them in the same order as the path $j \to 0$), then let i'' be the vertex immediately preceding j' on the path $i \to 0$. As above, by Lemma 25, $i'' \leqslant i$. Since i' is the first vertex encountered on the path $j \to 0$ with $i' \leqslant i$, and by construction this path encounters i'' before i', this implies that (i'', j') is not in the path $j \to 0$. Moreover, from Lemma 26 we get that $(i'', j') <_E (i, j)$, which shows that (i, j) is not externally active. We have thus shown that if $\mathcal{T}$ is a tree-tableau, then $S = S(\mathcal{T})$ is a spanning tree of G with $e(S) = 0$.

For the converse, suppose that $S = S(\mathcal{T})$ is a spanning tree with $e(S) = 0$. We show that $\mathcal{T}$ is a tree-tableau. Since $<_E$ is F-compatible, it follows that the edge e^ρ corresponding to the top left corner of F is minimal in $<_E$. Therefore it must be in the

spanning tree S, since otherwise it would be externally active. Thus the top left corner of $\mathcal{T}$ is dotted. Now consider a dotted cell e_1 of $\mathcal{T}$, and suppose it has both a dot above it in its column, say some e_2, and a dotted cell to its left in its row, say some e_3. Consider the cell e_4 that is in the same row as e_2 and column as e_3. Since e_1, e_2, e_4, e_3 is a cycle of F, and S is a spanning tree, e_4 cannot be in S. But the condition that $<_E$ is F-compatible implies that $e_4 <_E e_2 <_E e_1$ and $e_4 <_E e_3 <_E e_1$, and thus e_4 is externally active, hence the desired contradiction.

It remains to show that every dotted cell e has a dotted cell either above it in its column or to its left in its row. Let $e = (i, j)$ where i is the label of its row and j of its column. Since S is a spanning tree, there are paths $i \to 0$ and $j \to 0$ using edges of S, and exactly one of these paths starts along the edge e. Suppose this is the path $i \to 0$. We write this path as $i, e, j, e', i', \ldots, 0$. It is sufficient to show that the labels i, i' satisfy $i' < i$ since then the cell $e' = (i', j)$ is a dotted cell above e in its column. Seeking contradiction, suppose that $i' > i$. Let $e'' = (i'', j'')$ be the first edge on the path $i \to 0$ such that $i'' < i$ or $j'' < j$. By construction exactly one of these conditions will hold. If we have $j'' < j$, then the edge (i, j'') is externally active, since it is minimal in the cycle $i, j, i', \ldots, i'', j'', i$. Similarly, if $i'' < i$ then the edge (i'', j) is externally active. This is the desired contradiction, and thus $i' < i$ as desired. We can show in exactly the same fashion that if it is the path $j \to 0$ that goes through e, then there is a dotted cell to the left of e in its row. □

Thus, fixing an order $<_E$ that is F-compatible, we have a bijection between tree-tableaux and spanning trees with external activity equal to 0. We now use the work of Cori and Le Borgne [5] to build a new bijection from tree-tableaux to EW-tableaux.

7.3 Tree-tableaux and EW-tableaux

Let $G = G(F)$ be a Ferrers graph on $n + 1$ vertices, labeled $0, \ldots, n$. We present an algorithm that constructs an n-permutation $\tau = v_1 \ldots v_n$ from a tree-tableau $\mathcal{T}$ on the Ferrers diagram corresponding to G. This allows us to construct a bijection between tree-tableaux and EW-tableaux, which is different from that of Section 6.

Algorithm 27. Given a tree-tableau $\mathcal{T}$ with row and column labels $0, \ldots, n$, initialize $v_0 = 0$ and $i = 1$, and determine each v_i for $i > 0$ as follows:

1. Set $j = i - 1$.
2. If v_j is a column (resp. row) let e_i be the bottommost (resp. rightmost) dotted cell in v_j not in $\{e_1, \ldots, e_{i-1}\}$. If no such cell exists decrease j by one and repeat Step 2.
3. Let v_i be the row or column of e_i not in $\{v_0, \ldots, v_{i-1}\}$.
4. Increase i by one, if $i = n + 1$ then stop, otherwise go to Step 1.

The algorithm above can be described in the following way: Let e_1 be the rightmost dotted cell in the top row, set v_1 as the column of e_1, and say that this cell and column have been visited. We then take the bottommost dotted cell in column v_1, we set this

as e_2 and v_2 as the row of e_2. We then repeat this with the rightmost dotted cell in v_2. We continue this until there are no unvisited dotted cells in the current row/column v_i, at which point we go backwards through the visited rows and columns $v_{i-1}, \ldots, v_1$ until we find the first row/column that does have an unvisited dotted cell. We then take the rightmost/bottommost such cell, and then repeat the previously described process until all rows and columns have been visited.

Example 28. Let $\mathcal{T}$ be the tree-tableau in Figure 7.1, with rows and columns labeled. We construct the permutation $\tau = v_1 \ldots v_n$ by Algorithm 27. Start with $v_0 = 0$, which is the top row, and set $i = 1$ and $j = 0$. Since $v_j = 0$ is a row and $i = 1$, we set e_1 to be the rightmost dotted cell in v_0, that is, $e_1 = (0, 3)$. We then set $v_1 = 3$, increase i by one, so $i = 2$, and return to Step 1, setting $j = 1$. Now $v_j = v_1 = 3$, which is a column, so we let e_2 be the bottommost dotted cell in the column 3 that is not in $\{e_1\} = \{(0, 3)\}$. Thus $e_2 = (1, 3)$, $v_2 = 1$, and we set $i = 3$ and return to Step 1. Now we set $j = 2$, so $v_j = v_2 = 1$, which is a row. Thus e_3 is the rightmost dotted cell in the row 1 that isn't in $\{e_1, e_2\}$ so $e_3 = (1, 2)$ and $v_3 = 2$. We set $i = 4$ and return to Step 1.

Now $j = 3$, so $v_j = v_3 = 2$. But the column 2 contains a single dotted cell (which is e_3), so we must decrease j by one, that is, set $j = 2$ and $v_j = v_2 = 1$ and repeat Step 2. But the two dotted cells of the row 1 are e_2 and e_3, so again we decrease j by one, set $v_j = v_1 = 3$ and repeat Step 2. Once again, the two dotted cells e_1 and e_2 of the column 3 have already appeared, so again we decrease j by one, setting $v_j = v_0 = 0$. This time, the dotted cell $(0, 5)$ isn't in $\{e_1, e_2, e_3\}$, and it is the rightmost such dotted cell in the row 0, so we set $e_4 = (0, 5)$, $v_4 = 5$, $i = 5$ and return to Step 1, setting $j = 4$. This time $v_j = v_4 = 5$, and e_5 is the bottommost dotted cell in the column 5 that is not in $\{e_1, \ldots, e_4\}$, so $e_5 = (4, 5)$, $v_5 = 4$, and we increase i by one to $i = 6$. Now we have reached $i = n + 1$, so the algorithm terminates, producing the permutation $\tau = 31254$.

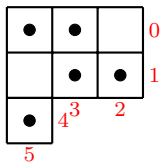

Figure 7.1: An example of a tree-tableau from which Algorithm 27 constructs the permutation $\tau = 31254$.

Lemma 29. *Algorithm 27 is well defined, and produces an n-permutation $\tau := v_1 \ldots v_n$.*

Proof. We need to show that for any $i < n$ there exists a vertex v_j for some $j \leqslant i$ such that the row or column v_j contains a dotted cell that isn't in $\{e_1, \ldots, e_{i-1}\}$. We proceed by contradiction. Fix $i < n$ such that for all $j \leqslant i$, all the dotted cells of the row or column v_j are in $\{e_1, \ldots, e_{i-1}\}$. Since $i < n$, there exists some dotted cell e in $\mathcal{T}$ that is not in $\{e_1, \ldots, e_{i-1}\}$. Without loss of generality, we may assume that e has a dotted cell to its left (and none above). Let v be the row containing e, and let e' be the leftmost dotted cell in row v. Since v can't be in $\{v_0, \ldots, v_i\}$, this means that e' is not in $\{e_1, \ldots, e_{i-1}\}$.

By definition, e' has no dotted cell to its left, so it must have one above. Let e'' be the topmost dotted cell in the column containing e'. As before, e'' cannot be in $\{e_1, \ldots, e_{i-1}\}$. By iterating these steps, we eventually reach a dotted cell f in the top row that isn't in $\{e_1, \ldots, e_{i-1}\}$. But this is a contradiction, since the top row is v_0, and by assumption all dotted cells in v_0 are in $\{e_1, \ldots, e_{i-1}\}$. □

Proposition 30. *Let $\tau = v_1 \ldots v_n$ be the n-permutation produced from a tree-tableau $\mathcal{T}$ by Algorithm 27. Then the descent bottoms of τ are exactly the row labels of the Ferrers diagram F of $\mathcal{T}$.*

Proof. We show by induction on $i \geqslant 1$ that

$$v_i \text{ is a descent of } \tau \text{ if and only if it is the label of a row of } F. \tag{7.1}$$

For $i = 1$, by construction e_1 is the rightmost dotted cell in the top row v_0, and v_1 is the column containing e_1. Since v_1 is not a descent of τ, this shows (7.1) for $i = 1$. Now fix $i \geqslant 1$, and suppose that we have shown (7.1) to hold for all $j \leqslant i$. We say a dotted cell e is *included* if it is in $\{e_1, \ldots, e_i\}$. We distinguish three cases.

Case 1 Suppose v_i is a column, and there are some non-included dotted cells in v_i. Then e_{i+1} is the bottommost of these dotted cells, and v_{i+1} the row containing e_{i+1}. Thus e_{i+1} is a cell of the Ferrers diagram F in row v_{i+1} and column v_i, so that by construction of the edge labels $v_{i+1} < v_i$, and v_{i+1} is a descent, as desired.

Case 2 Suppose v_i is a row, and there are some non-included dotted cells in v_i. As above, e_{i+1} is the rightmost of these cells, and it is a cell of F in row v_i and column v_{i+1} so that $v_{i+1} > v_i$, and v_{i+1} isn't a descent.

Case 3 Suppose all dotted cells of the row or column v_i are included. Let e^ρ be the dotted cell in the top left corner of $\mathcal{T}$. Given any dotted cell e, we can construct an "edge path" $e \to e^\rho$ by going from e to the dotted cell immediately above it in its column or immediately to its left in its row (by definition of a tree-tableau exactly one of these exists), and iterating until we reach e^ρ. We claim that the cell e_{i+1} must be the first dotted cell on the path $e_i \to e^\rho$ that is not in $\{e_1, \ldots, e_{i-1}\}$. If e_i is the first dot to be encountered with no cell below or to the right, then $e_i, \ldots, e_1$ forms the path $e_i \to e^\rho$ by definition. So e_{i+1} is the first non-included dotted cell in $e_i \to e^\rho$.

By induction suppose e_i is the k-th cell to be encountered with no dots below and right and let e_j be the $(k-1)$-th such cell, and assume the claim is true for e_j. By construction $v_j, \ldots, v_i$ forms a path from e_j to e_i, because in between these no other dotted cell is encountered with no dot below or to its right, so at each step we take the bottom or rightmost cell. Therefore, if the first non-included dotted cell we encounter, denoted v_t, is between v_j and v_i we take the bottommost or rightmost non-included dotted cell of v_t, which must be the first cell above or to the left of the previously included dotted cell of v_t, which implies e_{i+1} is the first dot on the path from e_i to e^ρ. If $t < j$, then by induction e_{i+1} is the first non-included dotted cell on the path $e_j \to e^\rho$, and the

path $e_i \to e^\rho$ is the concatenation of $e_i \to e_j$, where we didn't encounter a non-included dotted cell, and $e_j \to e^\rho$. So e_{i+1} must be the non-included dotted cell on the path $e_i \to e^\rho$.

Thus e_{i+1} is on the path $e_i \to e^\rho$, which implies e_{i+1} must be weakly northwest of e_i. Therefore, if v_{i+1} is a column, then the label of v_{i+1} is left of the label of v_i so $v_{i+1} > v_i$, and so v_{i+1} is not a descent bottom. If v_{i+1} is a row, then the label of v_{i+1} is above the label of v_i so $v_{i+1} < v_i$, so v_{i+1} is a descent bottom. This completes Case 3, and the proof. □

Theorem 31. *Let F be a Ferrers diagram and let $M : \mathrm{TLT}(F) \longrightarrow \mathrm{EWT}(F)$ be the map taking $\mathcal{T}$ to $\Psi^{-1}(\tau(\mathcal{T}))$, where $\tau(\mathcal{T})$ is the permutation constructed from $\mathcal{T}$ by Algorithm 27 and Ψ is the map from Definition 7. Then M is a bijection from tree-tableaux of shape F to* EW*-tableaux of shape F.*

Proof. Proposition 30, coupled with Proposition 11, shows that the map M is well defined. To show that M is a bijection, it is sufficient to show that it is injective. In fact, it suffices to show that $\mathcal{T} \mapsto \tau(\mathcal{T})$ is injective, since Ψ is a bijection by Proposition 12. Let $\mathcal{T}, \mathcal{T}'$ be two tree-tableaux and write $v_0, \ldots, v_n$ (resp. $v'_0, \ldots, v'_n$) and $e_1, \ldots, e_n$ (resp. $e'_1, \ldots, e'_n$) for the order in which vertices and edges appear in the construction of τ (resp. τ'). The construction implies that for any $i \geqslant 1$, $(v_0, \ldots, v_i) = (v'_0, \ldots, v'_i)$ if and only if $(e_1, \ldots, e_i) = (e'_1, \ldots, e'_i)$. Now $\mathcal{T} \neq \mathcal{T}'$ means that the sets of dotted cells of the two tableaux are different, and in particular $(e_1, \ldots, e_n) \neq (e'_1, \ldots, e'_n)$. This implies that $\tau(\mathcal{T}) = v_1 \cdots v_n \neq v'_1 \cdots v'_n = \tau(\mathcal{T}')$, as desired. □

Example 32. Let $\mathcal{T}$ be the tree-tableau on the left in Figure 7.2, with rows and columns labeled. This is the same tree-tableau as in Example 28 and in Figure 7.1, and the corresponding permutation is $\tau = 31254$. The EW-tableau on the right is constructed by first filling the top row with 1s, then reading τ from left to right, and filling in the remaining cells of v_i with 1s if v_i is the label of a row, and with 0s if it is the label of a column.

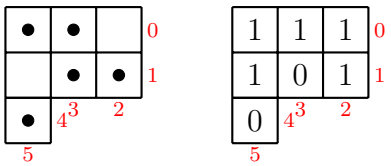

Figure 7.2: An example of a tree-tableau $\mathcal{T}$ (left) and its corresponding EW-tableau $\mathcal{T}'$ (right) via the bijection of Theorem 31, both corresponding to the permutation $\tau = 31254$.

Remark 33. The map in Theorem 31 can be combined with the results of Theorems 22 and 24 to obtain a bijection between spanning trees of the Ferrers graph $G = G(F)$ and minimal recurrent configurations of the ASM on G. In fact, this bijection is equivalent to, but presented in a different form from, the bijection obtained by composing those of Lemmas 2 and 4 in [5], restricted to spanning trees with no external activity, which map bijectively to minimal recurrent configurations. Indeed, it is possible to show that the

construction in Algorithm 27 yields the following property. Let $<_E$ be any F-compatible total order on the edges of G. Then for all $i \geqslant 0$, e_{i+1} is the largest edge, according to $<_E$, where the corresponding cell in $\mathcal{T}$ lies in exactly one column or row of $\{v_0, \ldots, v_i\}$, that is,

$$e_{i+1} = \max_{<_E}\{(w, w') \in S : w \notin \{v_0, \ldots, v_i\}, w' \in \{v_0, \ldots, v_i\}\}.$$

Remark 34. The bijection from Theorem 31 is different from the bijection in Section 6. Consider the tree-tableau on the left in Figure 7.3. The permutation τ obtained by Lemma 29 is $\tau = 4231$, which leads to the EW-tableau in the middle of Figure 7.3. However, the EW-tableau obtained through the bijection of Section 6 is the one on the right, and these are indeed different EW-tableaux.

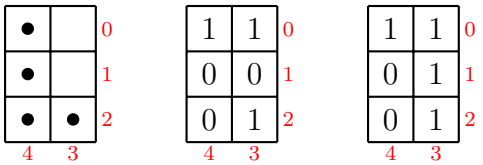

Figure 7.3: An example of a tree-tableau $\mathcal{T}$ (left), its corresponding EW-tableau $\mathcal{T}'$ and permutation via the bijection of Theorem 31 (center), and the corresponding EW-tableau and permutation via the bijection of Section 6. The two EW-tableaux are different.

Question 35. Although the EW-tableaux in Figure 7.3 are different, their corresponding permutations have a common prefix. In Example 32 we get the same EW-tableaux and thus the same permutation for both bijections, namely 31254. It is easy to see that the first letter of the corresponding permutation is always the same for both bijections; the question is what determines the maximal common prefix. In particular, can we find conditions under which the two bijections give the same EW-tableaux (and permutations)?

8 Statistics and properties of EW- and NEW-tableaux and their corresponding permutations

In this section we collect various results on EW- and NEW-tableaux, and their corresponding permutations, and statistics on these. We also mention two open problems and one conjecture that we have been unable to resolve.

Proposition 36. *There is a one-to-one correspondence between all-1 columns in a* EW- *or* NEW*-tableau* $\mathcal{T}$ *and fixed points in the permutation* $\mathrm{DE}(\Psi(\mathcal{T}))$.

Proof. Given an all-1 column C, we need to show that when constructing the permutation $\pi = \Psi(\mathcal{T})$ the label ℓ of column C is read after all rows and columns with smaller labels have been read and that the label read just before ℓ is smaller than ℓ. These two

conditions imply, according to the definition of the map DE in Proposition 6, that ℓ will be in place ℓ in $\mathrm{DE}(\pi)$.

The (label of a) column C is read only after all its 1s have been deleted, so if C is an all-1 column, all the rows with smaller labels will have been read before C is read. When that has been done, since columns are read from right to left, next all unread columns to the right of C will be read, since they will have been cleared of 1s, just as C. Thus all smaller labels will be read before ℓ, and one of them will be the last one read before ℓ.

Conversely, if ℓ is a fixed point of $\mathrm{DE}(\Psi(\mathcal{T}))$, then the label read just before ℓ in $\pi = \Psi(\mathcal{T})$ is smaller than ℓ, so that ℓ is the label of a column C. Moreover, all row labels that are smaller than ℓ are read before ℓ in π, so that by Corollary 14 the column C is all 1s. □

A *decreasing adjacency* in a permutation is a pair of adjacent letters $a_i a_{i+1}$ such that $a_i = a_{i+1} + 1$. For example, 65478132 has three decreasing adjacencies, 65, 54 and 32.

Lemma 37. *The number of fixed points different from 1 has the same distribution on n-permutations as does the number of decreasing adjacencies.*

Proof. The bijection DE translates a decreasing adjacency to a "minimal excedance", that is, an excedance whose top is one greater than its bottom. A cyclic right shift of the resulting permutation π takes each such minimal excedance to a fixed point. All fixed points in π arise this way, except 1, if 1 is a fixed point. □

Proposition 38. *The distribution of the number of all-zero-rows in* EW*-tableaux of size n is the same as the distribution of fixed points different from 1 in n-permutations, and the same as the distribution of decreasing adjacencies on such permutations.*

Proof. An all-0 row in an EW-tableau $\mathcal{T}$ corresponds to an all-1 column in the NEW-tableau $\mathcal{S} = \mathrm{Tr}(\mathcal{T})$, where Tr is the diagonal reflection of $\mathcal{T}$ described before Definition 16. The all-1 columns in $\mathcal{S}$ correspond one-to-one to fixed points in $\mathrm{DE}(\Psi(\mathcal{S}))$ except that 1 is a fixed point if and only if the top row of $\mathcal{S}$, and no other row, is all 1s. The latter part of the claim therefore follows from Lemma 37. □

Proposition 39. *The number of vertices in a longest directed path in a unique-sink acyclic orientation of a Ferrers graph $\mathcal{F}$ corresponding to an* EW*-tableau $\mathcal{T}$ equals the number of blocks in the run-decomposition of $\pi = \Psi(\mathcal{T})$. Moreover, any directed path in such an orientation contains at most one vertex label from each such block in π.*

Proof. As pointed out just after the proof of Lemma 5, an EW-tableau $\mathcal{T}$ corresponds uniquely to an acyclic orientation (with a unique sink) of the corresponding Ferrers graph $\mathcal{F}$. More precisely, a 1 in column c and row r corresponds to an edge from the vertex labeled c to the vertex labeled r in $\mathcal{F}$, a 0 to that edge being oriented the other way. A directed path through $\mathcal{F}$ must alternate between "top" and "bottom" vertices in the bipartite graph $\mathcal{F}$, corresponding, respectively, to row and column labels of $\mathcal{T}$, and

the contents of the cells in $\mathcal{T}$ corresponding to successive row and column labels must alternate between 0 and 1.

It follows from Lemma 10 that each block in the run-decomposition of $\pi = \Psi(\mathcal{T})$ contains only row labels or only column labels, and so a directed path in $\mathcal{F}$ must alternate between column and row label blocks in π. We claim that if there is an edge from a vertex u to a vertex v then v must precede u in π, which implies that a path can have at most one vertex from each block. Indeed, if u labels a row and there is an edge from u to v, then v is a column and the entry in the cell (u, v) is 0, so that by Corollary 14 v is read before u in π. An analogous argument shows that if u is a column label then the row label v must have been read before u.

We claim that the sequence consisting of the rightmost letter in each block in the run-decomposition of π gives the labels of vertices constituting a directed path in $\mathcal{F}$, when we read that sequence backwards, which will complete the proof.

When a column label is read, a 0 is cleared from all rows with a cell in that column whose labels are yet to be read; when a row label is read, a 1 is cleared from all columns with a cell in that row whose labels are yet to be read. The rightmost label in a block of column labels is the largest of them; the rightmost label in a block of row labels is the smallest of them. Thus, if c is the last label of a column in a block and r the last label of a row in the following block, then $c > r$, which implies that there is a cell in column c and row r, and that cell must have a 0, since c is read before r, implying that there is an edge from vertex r to vertex c in $\mathcal{F}$. An analogous argument shows that $r < c'$, where c' is the rightmost label in next block after that of r, and that there is a cell with a 1 in column c' and row r. □

Proposition 40. *The number of* EW*-tableaux of size n with exactly one 0 in each non-top row is the Fibonacci number F_{2n-2}. More precisely, the number of such tableaux whose first row has length k is $\binom{2n-1-k}{k-1}$.*

Proof. Call an EW-tableau 0-minimal if it has just one zero in each non-top row. Such a tableau is completely determined by the letters in each maximal set of adjacent cells at the bottom of columns in the same row, at most one of which letters can be a 0, and exactly one of which must be 0 in the case of the bottom row. This is because the cells directly above such a maximal set of cells can only have a 0 if there is a 0 in the bottom cell of the same column, since otherwise we would violate condition 3 in Definition 1. Therefore, every cell above a 1 must be filled with a 1 and every cell above a 0 must be filled with a 0 until a row that already contains a 0 is encountered and then everything from that row up must contain a 1.

Given a filling of each of the sets of maximal adjacent cells mentioned above, with at most one 0 in each set and exactly one 0 in the bottom row, we construct a binary string determined by the orientations of the edges on the NE-SW border of the tableaux and the content of their adjacent cells, starting from the edge just after the vertical edge of the top row. If the first edge is horizontal we write 01 for each edge in the maximal initial sequence of horizontal edges. For any other set of adjacent horizontal edges we write, starting from the right, 10 for an edge if there is a 1 in the cell above it, until we see

a 0 (if at all), at which point we write 01. For each of the remaining edges in the same maximal set of edges (that are at the same level) we write 01. For each vertical edge we write 00. Note that we read the edges of the tableau from right to left, but construct the string from left to right.

It is easy to see that the binary string thus constructed begins with a 0, ends with 01 and has no adjacent 1s. It is also easy to verify the following. Given a binary string w of length $2n-3$ with no adjacent 1s, let $\widehat{w} = 0w01$. Then $\widehat{w}$ is the binary string, of length $2n$, corresponding, as described above, to a unique EW-tableau of size n with a single 0 in each non-top row. Moreover, if $\widehat{w}$ has k occurrences of 1 (so w has $k-1$ occurrences of 1) then $\widehat{w}$ corresponds to such a tableau with top row of length k. It is easy to show that the number of binary strings of length $2n-3$ with $k-1$ occurrences of 1, none of them adjacent, is $\binom{2n-1-k}{k-1}$, and it is well known that the sum of those binomial coefficients, over all integers k, is the Fibonacci number F_{2n-2}, if $F_0 = F_1 = 1$. □

A permutation π is said to *contain the pattern* p, where p is also a permutation, if π has a subsequence of letters that appear in the same order of size as the letters of p, and π is said to *avoid* p if π has no such occurrence. For example, $\pi = 351624$ contains the pattern 231 in the subsequence 562, but π avoids 321, as it has no decreasing subsequence of length 3.

Proposition 41. *Let $\mathcal{T}$ be an* EW*-tableau with at least two rows, and $\pi = \Psi(\mathcal{T})$ the corresponding permutation. If π avoids* 231*, then the second row of $\mathcal{T}$ has a unique* 0*.*

Proof. If $\mathcal{T}$ has a second row, we first show that any column with a 0 in the second row of $\mathcal{T}$ must have all entries equal to 0 except in the top row. Let i be the label of the second row, and j the label of a column with a 0 in the second row. Then j appears before i in the permutation π by construction. Since $j > i$ and π avoids 231, this implies in particular that there cannot be any row label i' with $i' < j$ that appears before j in π. Thus, all entries in the column labeled j (other than the top row) must be a 0, since these correspond to such labels i', which must appear after j in π.

Suppose now that there are two columns with labels $j < j'$ whose entries in the second row are both 0. By the above, both columns must have all non-top entries equal to 0. By construction, such columns are read in increasing order at the start of the permutation π. In particular, j appears before j', and both appear before i (since their entries in the row labeled i are 0). But since $i < j < j'$, it follows that j, j' and i appearing in that order in π form a 231 pattern, which is forbidden. □

Question 42. Proposition 41 is not an equivalence; there are EW-tableaux with a unique 0 in the second row whose corresponding permutations do not avoid 231, such as the 2×2 tableau with a single 0, in the bottom right corner, whose corresponding permutation is in fact 231 itself. Is there some nice further condition on EW-tableaux that would make this an equivalence?

Proposition 43. *Let $\mathcal{T}$ be an* EW*-tableau, and $\pi = \Psi(\mathcal{T})$ the corresponding permutation. Then the following are equivalent.*

1. *The permutation π avoids the pattern* 213*.*

2. - *All rows of $\mathcal{T}$ are of the form $0\cdots01\cdots1$, where the sequence of* 1*s may be empty, and*
   - *in any row of $\mathcal{T}$, the leftmost* 1 *has no* 1 *in its column below it.*

Proof. We first show that (1) implies (2). Suppose π avoids 213. First, fix some row of $\mathcal{T}$, labeled i. Let j be the label of the column of the leftmost 1 in that row (if it exists). By construction we have $j > i$ and since the entry (i, j) is a 1, i appears before j in π. Let j' be the label of any column to the right of the column labeled j. Since $i < j' < j$ and π avoids 213, j' must also appear after i in π, and thus the entry (i, j') must be a 1. This shows the first point of (2).

For the second point, we need to show that all entries in the column labeled j that are below the row labeled i are 0. Let i' be the label of the row of such an entry. By construction we have $i < i' < j$. Moreover, recall that i appears before j in π. Since π avoids 213, i' must appear after i in π. Now by the above, the rows labeled i and i' are both of the form $0\cdots01\cdots1$. By construction, if the leftmost 1 in the row i' was to the left (or directly below) the leftmost 1 in the row labeled i, then i' would appear before i in π, which is not the case. In particular, the entry directly below the leftmost 1 in the row labeled i must be a 0, as desired.

We now show that (2) $\Rightarrow$ (1). Suppose that $\mathcal{T}$ satisfies both points of (2). We first show the following lemma.

Lemma 44. *If π contains a* 213*-pattern, then it contains a* 213*-pattern where the letter corresponding to the* 1 *of the pattern is the label of a row of $\mathcal{T}$.*

Proof. Let b, a, c be the letters of an occurrence of 213 corresponding, respectively, to $2, 1, 3$. Suppose a is the label of a column (otherwise we are done). The b cannot be the label of a column in the same block in the run-decomposition of π, since such blocks are increasing. Let a' be the leftmost letter in the block of a, so $a' \leqslant a$. By Lemma 15, the rightmost letter x of the block of row labels preceding a' is smaller than a', and thus b, x, c form the desired occurrence of 213. □

Seeking contradiction, suppose π contains an occurrence of 213, and write i, j, k for the letters corresponding respectively to $1, 2$ and 3. By Lemma 44, we may assume that i is the label of a row of $\mathcal{T}$. We distinguish three cases, where (i, j) refers to the cell in row i and column j of $\mathcal{T}$.

Case 1 Suppose j and k both label columns of $\mathcal{T}$. Since $i < j < k$, row i has entries in both those columns, and column j is to the right of column k. Moreover, since j, i, k appear in π in that order, the entry in (i, j) is a 0 and the entry in (i, k) is a 1. This contradicts the fact that each row is of the form $0\ldots01\ldots1$.

Case 2 Suppose j labels a column and k labels a row of $\mathcal{T}$. Since $i < j < k$, column j must have an entry in row i but not in row k. Since j appears before i in π, the entry in (i, j) is a 0. Because of the form of row i, all entries to the left of (i, j) must also be 0. In particular, all entries of row i in columns that also contain entries in row k are 0, which implies that k appears before i in π. This contradicts the assumption that j, i, k appear in π in that order.

Case 3 Suppose j labels a row of $\mathcal{T}$. Because $i < j$, row j is below row i in $\mathcal{T}$. Now the conditions of (2), together with the fact that j appears before i in π, imply that the entry of all the columns in row i that also have an entry in row j are 0s. This implies that k labels a row, since any column label greater than j must therefore appear before i in π. But $j < k$, so that all the columns in row i that have an entry in row j also have an entry in row k. Since these entries are all 0s, and i appears before k in π, this implies that the row k is all-0. But such a row would be read before i, since the column of the leftmost 1 of row i has no entry in row k. This is the desired contradiction. □

Definition 45. A row R in a tableau $\mathcal{T}$ *dominates* a row S in $\mathcal{T}$ if S lies above R and there is no column in $\mathcal{T}$ where R has a 0 and S has a 1. A tableau $\mathcal{T}$ is *domination-free* if no row of $\mathcal{T}$ dominates another.

Definition 46. A tableau $\mathcal{T}$ is *top-justified* if all the 1s in each column are at the top, that is, no 1 has a 0 above it in the same column.

Lemma 47. *A domination-free tableau $\mathcal{T}$ is top-justified.*

Proof. Note first that no two rows of $\mathcal{T}$ can have the same filling or the lower one would dominate the upper row. Given a row u above a row r in $\mathcal{T}$, the filling of the cells in r can of course be obtained from the filling of u by changing some entries. We cannot only change some 0s to 1s, since then r would dominate u, so we must change some 1s in u to 0s in r. If we then change some 0s in u to 1s in r we would get the kind of pattern forbidden by the definition of NEW-tableaux. Therefore, any cell containing a 1 must have only 1s above it. □

Recall the definition of a pattern in a permutation, given at the beginning of this section. A permutation π avoids the *vincular pattern* 32–1 if π has no decreasing subsequence of length 3 whose first two letters are adjacent. For example, 4132 avoids 32–1, whereas 4231 does not, because of the subsequence 421. Recall also that a *right-to-left minimum* of a permutation π is a letter of π that is smaller than all letters to its right. For example, 3172546 has right-to-left minima $1, 2, 4, 6$.

Proposition 48. *A domination-free* NEW*-tableau $\mathcal{N}$ of size n has exactly k columns containing a 0 if and only if the permutation $\pi = \Psi(\mathcal{N})$ begins with 1, avoids the vincular pattern* 32–1 *and has $(n-k)$ right-to-left minima. Thus, such* NEW*-tableaux are counted by $S(n-1, n-1-k)$, the Stirling number of the second kind, that is, the number of partitions of a set of $(n-1)$ elements into $(n-1-k)$ blocks.*

Proof. By Lemma 47, $\mathcal{N}$ is top-justified, so every column has a cell that contains a 1 with only 1s above and no 1s below, and we call such a cell a *transition cell* if it has some 0 below it (that is, is not the bottom cell in its column). This implies the top row is all 1s and that every non-bottom row must contain a transition cell, since a 1 must be lost in some column going from a row to the next one below. We let ℓ denote the number of rows and z the number of columns with no zeros, so $n - k = \ell + z$.

First we show the forward direction. Given a NEW-tableau $\mathcal{T}$ with no dominating rows, the first letter of $\pi = \Psi(\mathcal{T})$ must be a 1 because only the top row is all 1s. Next note that on a pass from bottom to top we will never find two rows that have no 0s, since otherwise the lower row would dominate the upper one. Also note that a row cannot be only 1s unless every row above is only 1s, which means all the rows above must already have been emptied. Moreover, if all rows above are empty, then all columns with a smaller label must be empty as well. So on each pass from bottom to top we encounter exactly one row that we can empty, which gives a letter that is a right-to-left minimum. As these are the only elements that can be the bottom of a descent this implies that π avoids 32–1. So every row contributes a right-to-left minimum. In addition, the label of an all-1 column appears in the block immediately after the label of its bottom row in the run-decomposition of π, and since column labels are read in increasing order, the all-1 columns contribute z extra right-to-left minima (see the end of the previous paragraph).

For the other direction, suppose $\pi = \Psi(\mathcal{T})$ avoids 32–1, has t right-to-left minima and begins with 1. Let s, r be two row labels of $\mathcal{T}$ such that row s is above row r. Then $r \neq 1$ so r is not the first letter of π and is therefore a descent bottom by point 2 of Theorem 17. Since π avoids 32–1, this implies in particular that r must appear after s in π. Let c be the column label immediately preceding r. By construction, $c > r > s$ and s, c, r appear in π in that order, so the column c has a 0 in row r and a 1 in row s, so that row r does not dominate row s. Thus $\mathcal{T}$ is domination-free. Moreover, we have shown that the labels of rows appear in π in order from top to bottom, so that each is a right-to-left minimum of π. As above, any all-1 column is also a right-to-left minimum. Finally, a column c containing a 0 in row r must be read before r in π, and since $c > r$, c is not a left-to-right minimum. Thus the number of columns containing a 0 is $n - t$ as desired.

Finally, the fact that the permutations described in the statement of the theorem are counted by the Stirling numbers of the second kind was shown in [4, Porism 1]. □

Proposition 49. *There is a bijection between top-justified* NEW*-tableaux of size n with k rows and set partitions of $[n] = \{1, \ldots, n\}$ with k blocks. The number of such* NEW*-tableaux thus equals $S(n, k)$, the Stirling number of the second kind.*

Proof. Let $\mathsf{N}(n, k)$ be the set of top-justified NEW-tableaux of size n with k rows, and $\mathsf{P}(n, k)$ the set of set partitions of $[n]$ with k blocks. Given $\mathcal{N} \in \mathsf{N}(n, k)$, we construct a set partition $F(\mathcal{N}) = \mathcal{P}_1, \ldots, \mathcal{P}_k \in \mathsf{P}(n, k)$ as follows. Let $r_1, \ldots, r_k$ be the row labels of $\mathcal{N}$. For all $j \in \{1, \ldots, k\}$, we define $\mathcal{P}_j$ to be the set consisting of r_j and of all column labels c such that the bottommost 1 of the column labeled c is in the row labeled r_j. We claim that $F : \mathsf{N}(n, k) \to \mathsf{P}(n, k)$ is a bijection.

Each column has exactly one bottommost 1, so that each column label belongs to

exactly one block of $\mathcal{P}$, as does each row label by construction. Thus $\mathcal{P}$ is a set partition of $[n]$ with k blocks, and F is well defined. To see that F is a bijection, we construct its inverse.

Let $\mathcal{P} = \mathcal{P}_1, \ldots, \mathcal{P}_k \in \mathsf{P}(n,k)$. For all $j \in \{1, \ldots, k\}$, let r_j be the minimum element of the block $\mathcal{P}_j$. Let $\mathcal{F}$ be the Ferrers diagram whose row labels are $r_1, \ldots, r_k$. We define a filling $\mathcal{N} = G(\mathcal{P})$ of $\mathcal{F}$ as follows. Any element $c \neq r_j$ in block $\mathcal{P}_j$ is greater than r_j and labels a column of $\mathcal{F}$, and thus row r_j has a cell in column c. Fill that cell and all the cells above it in column c with 1s, and all the cells below that cell in column c with a 0. This clearly defines a filling where every column has a 1 and no 1 has a 0 below it in the same column.

Moreover, no such filling of a Ferrers diagram can contain a rectangle violating condition 2 of Definition 16, since there cannot be a 0 above a 1 in a column. Thus $\mathcal{N} \in \mathsf{N}(n,k)$ as desired, and it is easy to see that the maps F and G are inverses of each other. □

Proposition 50. *If the permutation $\pi = \Psi(\mathcal{N})$ for a* NEW*-tableau $\mathcal{N}$ avoids the pattern 231 then $\mathcal{N}$ has no 0 above a 1 in a column.*

Proof. Suppose $\mathcal{N}$ has a column labeled k that has a 0 in a row labeled i above a 1 in a row labeled j. By construction $i < j < k$, and from point 3 of Theorem 17 the labels j, k, i appear in that order in π, which forms a 231 pattern. □

Propositions 49 and 50 suggest the following problem, since n-permutations avoiding 231 are equinumerous with non-crossing[3] set partitions of $[n]$ (both counted by the Catalan numbers):

Problem 51. Find a bijection from the NEW-tableaux in Proposition 49 to set partitions, such that the set partition corresponding to a NEW-tableau $\mathcal{N}$ is non-crossing if and only if the permutation $\pi = \Psi(\mathcal{N})$ avoids 231.

Note that the bijection F from the proof of Proposition 49 does not solve this problem. Indeed, consider the NEW-tableau $\mathcal{N}$ with two rows and two columns, whose entries read by rows are $[1,1], [1,0]$. We have $\mathcal{N} \in \mathsf{N}(4,2)$. However, the permutation $\pi = \Psi(\mathcal{N})$ is given by 1324, which avoids 231, whereas the set partition $\mathcal{P} = F(\mathcal{N})$ is given by $\{1,3\}, \{2,4\}$, which has a crossing.

A *big descent* in a permutation $\pi = a_1 a_2 \ldots a_n$ is an index i such that $a_i \geqslant a_{i+1} + 2$.

Proposition 52. *The number of* NEW*-tableaux of size n with k columns containing a 0 equals the number of n-permutations with k big descents.*

Proof. Proposition 38 implies that the distribution of the number of rows containing at least one 1 in EW-tableaux of size n is the same as the distribution of big descents on n-permutations, because decreasing adjacencies are precisely those descents that aren't big. But in the construction at the beginning of Section 3, NEW-tableaux of size n with k columns containing a 0 correspond exactly to EW-tableaux of size n with k rows containing a 1, which proves the claim. □

[3] A set partition has a crossing if it has a block containing a and c and another block containing b and d, where $a < b < c < d$.

It follows from Theorem 3 (and already from [14, Cor. 4.5]) that EW-tableaux of size n with $k+1$ rows are equinumerous with n-permutations with k excedances. It follows, of course, that the number of such tableaux with k columns equals the number of n-permutations with k non-excedances. We have however not been able to prove the following innocent-looking conjecture, which we have verified for $n \leqslant 12$.

Conjecture 53. The number of EW-tableaux of size n with k columns containing a 0 equals the Eulerian number $A(n,k)$ counting n-permutations with k excedances.

References

[1] The on-line encyclopedia of integer sequences, http://oeis.org/.

[2] J.-C. Aval, A. Boussicault, and P. Nadeau. Tree-like tableaux. *Electron. J. Combin.*, 20(4):#P4, 2013.

[3] J.-C. Aval, M. D'Adderio, M. Dukes, and Y. Le Borgne. Two operators on sandpile configurations, the sandpile model on the complete bipartite graph, and a cyclic lemma. *Adv. in Appl. Math.*, 73:59–98, 2016.

[4] A. Claesson. Generalized pattern avoidance. *European J. Combin.*, 22(7):961–971, 2001.

[5] R. Cori and Y. Le Borgne. The sand-pile model and Tutte polynomials. *Adv. Appl. Math.*, 30(1-2):44–52, 2003.

[6] S. Corteel, O. Mandelshtam, and L. Williams. Combinatorics of the two-species ASEP and Koornwinder moments. arXiv:1510.05023.

[7] S. Corteel and P. Nadeau. Permutation tableaux and permutation descents. *Europ. J. Combin*, 30:295–310, 2009.

[8] S. Corteel and L. K. Williams. A Markov chain on permutations which projects to the PASEP. *Int. Math. Res. Not. IMRN*, (17):Art. ID rnm055, 27, 2007.

[9] S. Corteel and L. K. Williams. Tableaux combinatorics for the asymmetric exclusion process. *Adv. in Appl. Math.*, 39(3):293–310, 2007.

[10] D. Dhar. Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.*, 64:1613–1616, Apr 1990.

[11] M. Dukes and Y. Le Borgne. Parallelogram polyominoes, the sandpile model on a complete bipartite graph, and a q,t-Narayana polynomial. *J. Combin. Theory Ser. A*, 120(4):816–842, 2013.

[12] M. Dukes, T. Selig, J. P. Smith, and E. Steingrímsson. EW-tableaux, permutations and recurrent configurations of the Abelian sandpile model. In preparation.

[13] R. Ehrenborg and E. Steingrímsson. The excedance set of a permutation. *Adv. in Appl. Math.*, 24(3):284–299, 2000.

[14] R. Ehrenborg and S. van Willigenburg. Enumerative properties of Ferrers graphs. *Discrete Comput. Geom.*, 32(4):481–492, 2004.

[15] D. Foata and M.-P. Schützenberger. *Théorie géométrique des polynômes eulériens*. Lecture Notes in Mathematics, Vol. 138. Springer-Verlag, Berlin-New York, 1970.

[16] C. Greene and T. Zaslavsky. On the interpretation of Whitney numbers through arrangements of hyperplanes, zonotopes, non-Radon partitions, and orientations of graphs. *Trans. Amer. Math. Soc.*, 280(1):97–126, 1983.

[17] A. Postnikov. Total positivity, grassmannians, and networks. arXiv:math/0609764.

[18] M. Schulz. Minimal recurrent configurations of chip firing games and directed acyclic graphs. In N. Fatès, J. Kari, and T. Worsch, editors, *Automata*, Discrete Mathematics and Theoretical Computer Science Proceedings, pages 111–124. DMTCS, 2010.

[19] E. Steingrímsson. *Permutations statistics of indexed and poset permutations*. ProQuest LLC, Ann Arbor, MI, 1992. Thesis (Ph.D.)–Massachusetts Institute of Technology.

[20] E. Steingrímsson. Permutation statistics of indexed permutations. *European J. Combin.*, 15(2):187–205, 1994.

[21] E. Steingrímsson and L. K. Williams. Permutation tableaux and permutation patterns. *J. Combin. Theory Ser. A*, 114(2):211–234, 2007.

[22] X. Viennot. Alternative tableaux, permutations and partially asymmetric exclusion process, Isaac Newton Institute, April 2007, video and slides from a talk at the workshop *Statistical-Mechanics and Quantum-Field Theory Methods in Combinatorial Enumeration*, http://www.newton.ac.uk/webseminars/pg+ws/2008/csm/csmw04/0423/viennot.