Advances in Applied Mathematics 96 (2018) 195–215

Contents lists available at ScienceDirect

Advances in Applied Mathematics

www.elsevier.com/locate/yaama

Bijective enumerations of Γ-free 0–1 matrices

Beáta Bényi [a], Gábor V. Nagy [b,*]

[a] *Faculty of Water Sciences, National University of Public Service, Hungary*
[b] *Bolyai Institute, University of Szeged, Hungary*

ARTICLE INFO

Article history:
Received 2 August 2017
Received in revised form 6 December 2017
Accepted 12 December 2017

MSC:
05A05
05A19

Keywords:
Γ-free matrix
Non-ambiguous forest

ABSTRACT

We construct a new bijection between the set of $n \times k$ 0–1 matrices with no three 1's forming a Γ configuration and the set of (n, k)-Callan sequences, a simple structure counted by poly-Bernoulli numbers. We give two applications of this result: We derive the generating function of Γ-free matrices, and we give a new bijective proof for an elegant result of Aval et al. that states that the number of complete non-ambiguous forests with n leaves is equal to the number of pairs of permutations of $\{1, \ldots, n\}$ with no common rise.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

We call a 0–1 matrix Γ-*free* if it does not contain 1's in positions such that they form a Γ configuration; i.e. two 1's in the same row and a third 1 below the left of these in the same column. $\Gamma = \begin{smallmatrix} 1 & 1 \\ 1 & * \end{smallmatrix}$. For instance, matrix A is not Γ-free because the bold 1's form a Γ configuration, while matrix B is a Γ-free matrix.

$$A = \begin{pmatrix} 0 & \mathbf{1} & \mathbf{1} & 0 \\ 1 & 1 & 0 & 0 \\ 0 & \mathbf{1} & 0 & 1 \end{pmatrix}, \quad B = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 \end{pmatrix}$$

* Corresponding author.
E-mail addresses: beata.benyi@gmail.com (B. Bényi), ngaba@math.u-szeged.hu (G.V. Nagy).

https://doi.org/10.1016/j.aam.2017.12.002
0196-8858/© 2018 Elsevier Inc. All rights reserved.

Table 1
The poly-Bernoulli numbers $B_n^{(-k)}$.

n, k	0	1	2	3	4	5
0	1	1	1	1	1	1
1	1	2	4	8	16	32
2	1	4	14	46	146	454
3	1	8	46	230	1066	4718
4	1	16	146	1066	6906	41506
5	1	32	454	4718	41506	329462

Clearly, we can say that a matrix is Γ-free if and only if it does not contain any of the submatrices from the following set:

$$\left\{\begin{pmatrix}1 & 1\\ 1 & 0\end{pmatrix}, \begin{pmatrix}1 & 1\\ 1 & 1\end{pmatrix}\right\}.$$

Pattern avoidance is an important notion in combinatorics. Matrices were also investigated from different point of view in this context; both extremal [9] and enumerative [11], [12] results are known.

Γ-free 0–1 matrices of size $n \times k$ contain at most $n + k - 1$ 1's [9]. The set of $n \times k$ 0–1 Γ-free matrices is one of the matrix classes that are enumerated by the poly-Bernoulli numbers, $B_n^{(-k)}$ [4]. Besides matrix classes that are characterized by excluded submatrices there are several other combinatorial objects that are enumerated by the poly-Bernoulli numbers. For instance, permutations with a given exceedance set, permutations with a constraint on the distance of their values and images, Callan permutations, acyclic orientations of complete bipartite graphs, non-ambiguous forests, etc. For further details, including recurrence relations and the original definition of poly-Bernoulli numbers via generating function, see [4], [5] and [6]. There is also a nice combinatorial formula of the poly-Bernoulli numbers of negative k indices: For $k > 0$,

$$B_n^{(-k)} = \sum_{m=0}^{\min(n,k)} m! \left\{ {n+1 \atop m+1} \right\} m! \left\{ {k+1 \atop m+1} \right\}, \tag{1}$$

where $\left\{ {n \atop m} \right\}$ denotes a Stirling number of the second kind. Table 1 shows the values of $B_n^{(-k)}$ for small k and n.

From (1), we give an obvious combinatorial interpretation of the numbers $B_n^{(-k)}$, which will be regarded as their combinatorial definition in this paper. (This interpretation is essentially the same as the one that counts Callan permutations.) On an (n, k)-*Callan sequence* we mean a sequence $(S_1, T_1), \ldots, (S_m, T_m)$ for some $m \in \mathbb{N}_0$ such that $S_1, \ldots, S_m$ are pairwise disjoint nonempty subsets of $\{1, \ldots, n\}$, and $T_1, \ldots, T_m$ are pairwise disjoint nonempty subsets of $\{1, \ldots, k\}$. We note that the empty sequence is also a Callan sequence with $m = 0$.

Lemma 1. *For $k > 0$, $B_n^{(-k)}$ counts the number of (n, k)-Callan sequences.*

Proof. For a fixed length m, there are $m!\left\{ {n+1 \atop m+1} \right\}$ ways to give the sequence $S_1, \ldots, S_m$. This is because if we partition $\{1, \ldots, n+1\}$ into $m+1$ (nonempty) classes, and order the m classes not containing the element $n+1$ arbitrarily, then we obtain every possible $S_1, \ldots, S_m$ sequences uniquely in this way. ($n+1$ can be thought as a "dummy" element, which is introduced to identify the class containing the elements of $\{1, \ldots, n\} \setminus \cup_{i=1}^{m} S_i$ and to allow this set to be empty.) Analogously, there are $m!\left\{ {k+1 \atop m+1} \right\}$ ways to give the sequence $T_1, \ldots, T_m$. The sequences $(S_i)_{i=1}^{m}$ and $(T_i)_{i=1}^{m}$ are independent from each other and, trivially, m can be at most $\min(n, k)$, so the statement follows from formula (1). □

Now we can state our first main theorem.

Theorem 2. *There exists a bijection between the set of* Γ*-free* $n \times k$ 0-1 *matrices and the set of* (n, k)*-Callan sequences. Thus the number of* Γ*-free* $n \times k$ 0-1 *matrices is* $B_n^{(-k)}$.

As mentioned earlier, the Γ-free 0–1 matrices have already been enumerated in [4]. In that paper, Bényi and Hajnal take an obvious combinatorial interpretation of poly-Bernoulli numbers (which is basically the same as ours, Lemma 1) and give a bijective proof for the second statement of Theorem 2. Their proof involved a lot of technical details and it was still desirable to find a simple bijective explanation for this result; a direct simple bijection that exhibits the connection between Γ-free matrices and Callan sequences. In Section 2 we define such a bijection, which is essentially different from the one in [4]. It reveals the inner structure of Γ-free matrices from a new point of view. As an application, we derive the generating function of Γ-free matrices in Section 3. Furthermore, we use our bijection in Section 4 to prove bijectively that the set of complete non-ambiguous forests and the set of permutation pairs with no common rise are equinumerous.

Γ-free matrices are very closely related to non-ambiguous trees and forests that were introduced in [2]. A *non-ambiguous tree* of size n is a set A of n points: $v = (x(v), y(v)) \in \mathbb{N}^+ \times \mathbb{N}^+$ satisfying the following conditions:

1. $(1, 1) \in A$ is the root of the tree;
2. for a given non-root point $p \in A$, there exists one point $q \in A$ such that $y(q) < y(p)$ and $x(q) = x(p)$ or one point s such that $x(s) < x(p)$ and $y(s) = y(p)$ but not both;
3. there is no empty line between two given points: if there exists a point $p \in A$ such that $x(p) = x$ (resp. $y(p) = y$), then for every $x' < x$ (resp. $y' < y$) there exists a point $q \in A$ such that $x(q) = x'$ (resp. $y(q) = y'$).

This structure can be viewed as a rooted binary tree graph on vertex set A with root $(1, 1)$ in which the parent of a non-root vertex p is the nearest q or s from condition 2. The name "non-ambiguous" comes from the property that the parent of p can be uniquely recovered

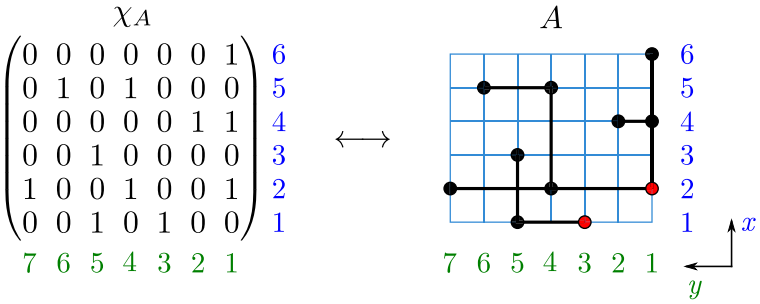

Fig. 1. A Γ-free 0–1 matrix and the corresponding non-ambiguous forest. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

from the vertex set A, since either q or s does not exist. In order to be consistent with the notion of Γ-free matrices, we slightly modified the original definition of non-ambiguous trees (we translated A), and we follow an unusual convention in this paper:

Convention 3. Throughout this paper, the rows of a matrix are always indexed *from bottom to top* and the columns are indexed *from right to left*.

The *characteristic matrix* of a finite set $A \subseteq \mathbb{N}^+ \times \mathbb{N}^+$ is a 0–1 matrix χ_A with $\max_{v\in A} x(v)$ rows and $\max_{v\in A} y(v)$ columns, such that, *following Convention 3*, there is a 1 in position (i, j) of χ_A, if and only if $(i, j) \in A$. Clearly, the characteristic matrix of a non-ambiguous tree is Γ-free by condition 2, so non-ambiguous trees can be thought as special Γ-free 0–1 matrices. It turns out that there is a graph theoretic terminology for Γ-free 0–1 matrices in literature. A *non-ambiguous forest* is a finite set $A \subseteq \mathbb{N}^+ \times \mathbb{N}^+$ such that χ_A is a Γ-free matrix without all-0 rows and columns. This is actually the original definition given in [2], but alternatively we could say that A is a non-ambiguous forest, if it satisfies condition 3 in the definition of non-ambiguous trees and the modified condition 2 that allows the possibility that neither such a q nor such an s exists. Analogously to trees, a non-ambiguous forest has an underlying (rooted) binary forest structure, see Fig. 1. We note that non-ambiguous trees are exactly those non-ambiguous forests that have one component. We say that a non-ambiguous tree or forest is *complete* if its vertices have either 0 or 2 children.

In [3] Aval et al. presented a bijection between non-ambiguous trees and the corresponding subclass of Callan permutations using associated ordered trees as intermediate combinatorial objects. As an easy corollary of the proof of Theorem 2, we deduce the number of non-ambiguous forests with characteristic matrix of size $n \times k$ in Section 2. A special class of non-ambiguous trees has interesting connection to the J-Bessel function. Let b_n denote the number of complete non-ambiguous trees with n internal vertices (in OEIS [13] as sequence A002190). Then [2]:

$$\sum_{n\geqslant 0} b_n \frac{x^{n+1}}{(n+1)!^2} = -\ln J_0(2\sqrt{x}),$$

where J_0 is the Bessel function. Aisbett [1] and Jin [10] show strong connections of complete non-ambiguous trees with a poset of vector partitions.

Let $\alpha = (a_1, \dots, a_n)$ and $\beta = (b_1, \dots, b_n)$ be permutations of $\{1, \dots, n\}$. We say the pair (α, β) has *common rise* at position i (where $1 \leqslant i \leqslant n-1$), if $a_i < a_{i+1}$ and $b_i < b_{i+1}$ hold at the same time. In [2] Aval et al. made a nice observation. As a corollary, they noted that the number of complete non-ambiguous forests with n leaves has the same generating function as the number of pairs of permutations of $\{1, \dots, n\}$ with no common rise [7], denoting these quantities by $\tau(n)$ and $\omega(n)$, respectively,

$$\sum_{n \geqslant 0} \tau(n) \frac{x^n}{n!^2} = \left(\sum_{n \geqslant 0} \frac{(-1)^n x^n}{n!^2} \right)^{-1} = \sum_{n \geqslant 0} \omega(n) \frac{x^n}{n!^2};$$

and asked for a bijective explanation of the equality $\tau(n) = \omega(n)$. Jin described a bijection with the intermediate combinatorial objects of certain heaps [10]. We present a more direct bijection in Section 4, based on a labeled forest structure. This is our second main result.

Theorem 4. *There exists a bijection between the set of complete non-ambiguous forests with n leaves and the set of pairs of permutations of $\{1, \dots, n\}$ with no common rise. Thus these sets are equinumerous.*

In fact, we prove this theorem in a stronger form; we construct a bijection from the set of complete non-ambiguous forests with a fixed set of leaves to the corresponding subset of pairs of permutations with no common rise; see Theorem 11 in Section 4.

2. The number of Γ-free 0–1 matrices

In this section we present a proof of Theorem 2. But we need to cite a folklore result first.

By *rooted forest* we mean a vertex-disjoint union of (unordered) rooted trees. Two rooted forests with the same vertex set are considered the same, if and only if they have the same set of roots and they have the same edge set. In our proofs, the edges will usually be directed from parent to child (so the components become arborescences). Fix a totally ordered set $(V, <)$. We say that a rooted forest F on vertex set V is *increasing*, if whenever the vertex u is the parent of vertex v in F, then $u < v$. (See Fig. 2 for an example.) We note that in our figures we always list the children of a given parent in decreasing order from left to right (their order "does not count"), the tree components are listed in the decreasing order of their roots, and we follow this order in our algorithms, too.

Notation 5. If v is a vertex of a rooted forest F, then $F[v]$ denotes the rooted subtree of F spanned by v and its descendants (children, grandchildren etc.). The root of $F[v]$ is v.

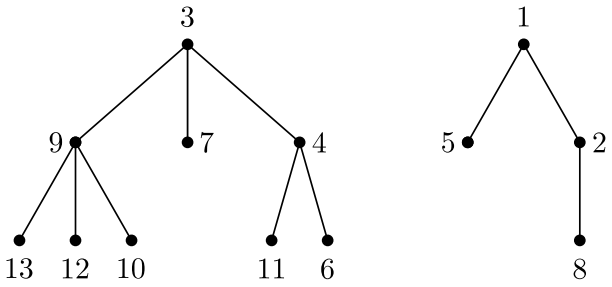

Fig. 2. An increasing forest on vertex set $\{1, \dots, 13\}$.

The following lemma is well known [14]:

Lemma 6. *Let V be a finite, totally ordered set. There exists a bijection π from the set of increasing forests on vertex set V to the set of permutations of V.*

Proof. We can obviously assume, that $V = \{1, \dots, n\}$ for some n, equipped with the natural order. Let F be an increasing forest on V. We will need the *pre-order transversal* of a rooted tree T, denoted by $\alpha(T)$. It is a permutation of the vertices of T which can be defined recursively as follows: Let r be the root of T. If T has only one vertex, r, then $\alpha(T) := r$; otherwise let $v_1, \dots, v_m$ be the children of r in decreasing order, and set $\alpha(T) := r\alpha(T[v_1]) \dots \alpha(T[v_m])$. (The "product" means concatenation here.) Now, if the (rooted tree) components of F are $T_1, \dots, T_k$, listed in the decreasing order of their roots, then $\pi(F)$ is defined to be the permutation $\alpha(T_1)\alpha(T_2) \dots \alpha(T_k)$. For example, π maps the increasing forest in Fig. 2 to the permutation $(3, 9, 13, 12, 10, 7, 4, 11, 6, 1, 5, 2, 8)$, for $n = 13$. It is straightforward to check that π is a bijection, as the lemma states. For further details (with a slightly different terminology), see Section 1.5 in [14]. □

Corollary 7. [14] *The number of increasing forests on vertex set $\{1, \dots, n\}$ is $n!$.*

Now we are ready to prove our first main theorem.

Proof of Theorem 2. Let $\mathcal{M}$ denote the set of Γ-free $n \times k$ 0-1 matrices, and let $\mathcal{S}$ denote the set of (n, k)-Callan sequences. Now we define a function $\phi\colon \mathcal{M} \to \mathcal{S}$, which will be proved to be bijective. An example is shown in Fig. 3; the details will be explained in Figs. 4–7.

Let M be an arbitrary matrix from $\mathcal{M}$. The rows of M can be identified with the numbers $1, \dots, n$, and the columns can be identified with the numbers $1, \dots, k$, and we index rows and columns as described in Convention 3.

We say that an element 1 in M is a *top*-1, if it is the highest 1 in its column, i.e. if there is no 1 above it in its column. There are three types of rows in M: There are the all-0 rows; there are the rows that contain at least one top-1, we call them *top rows*; and there are the rows that contain at least one 1 but none of their 1's is top-1, we call them

$$M = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{matrix} 11 \\ 10 \\ 9 \\ 8 \\ 7 \\ 6 \\ 5 \\ 4 \\ 3 \\ 2 \\ 1 \end{matrix}$$

$$\begin{matrix} 12 & 11 & 10 & 9 & 8 & 7 & 6 & 5 & 4 & 3 & 2 & 1 \end{matrix}$$

$\phi(M)$: $(\{2,7,11\},\{5,6,8\})$, $(\{8\},\{3\})$, $(\{6\},\{7\})$, $(\{9\},\{2,11\})$, $(\{4,5\},\{4,12\})$, $(\{1\},\{1\})$, $(\{3\},\{9,10\})$.

Fig. 3. A Γ-free matrix and the corresponding Callan sequence ($n = 11$, $k = 12$).

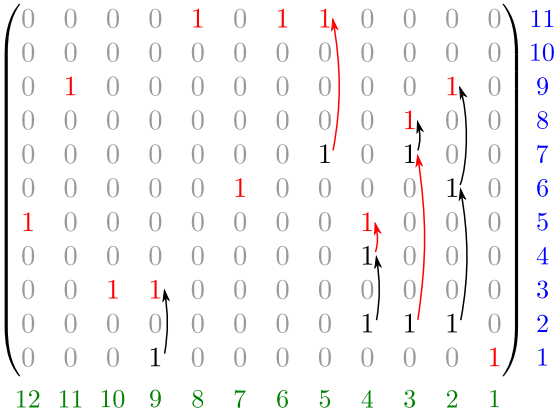

Fig. 4. The auxiliary graph G_M. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

special rows. In order to reduce the amount of required formalism, we define $\phi(M)$ with the help of auxiliary structures.

The auxiliary graph G_M. Let $G = G_M$ be the following directed graph: The vertices of G are the (positions of) 1's of M. For each non-top 1 vertex u, we add a directed edge starting from u and ending at the next (lowest) 1 above u in its column; and there are no other edges in G. If the edge e starts from row s and ends in row t, we say that the *length* of e is $t - s$. Since M is Γ-free, the edges starting from an arbitrary fixed row have pairwise distinct lengths. For each special row s, the longest edge starting from s is called *special*. We note that this is a valid definition and we underline that no special edge starts from a top row. The non-special edges of G are called *regular*. See Fig. 4 for an illustration, where the top-1's and the special edges are colored red.

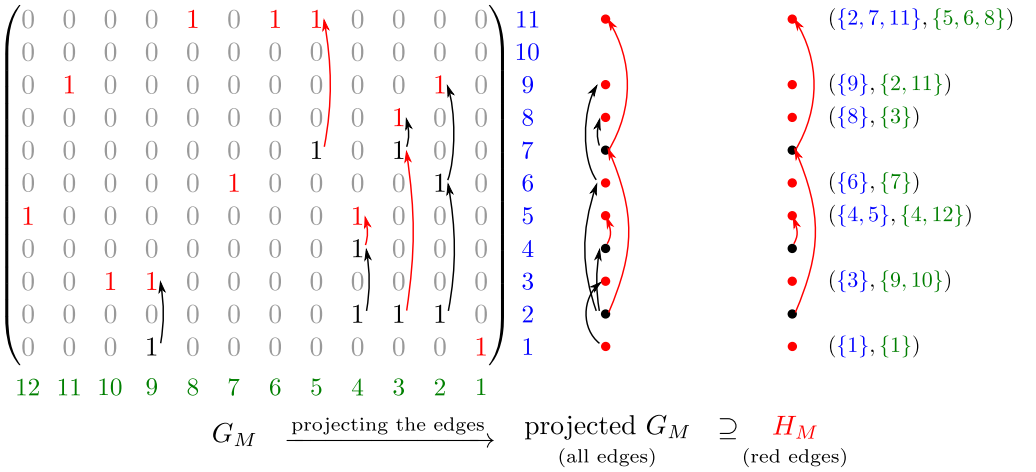

Fig. 5. The auxiliary graph H_M and the assigned pairs (R, C). (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

The auxiliary graph H_M. The next step is to "project the special edges horizontally": Let $H = H_M$ be the directed graph whose vertices are the *non-all-0 rows* (row indices) of M, and for each *special* edge e of G, there is an edge e' in H such that if e starts from row s and ends in row t, then e' starts from vertex s and ends at vertex t; and there are no other edges in H. H has a very simple structure. Since all edges are directed "upwards", there is no directed cycle in H. All vertices corresponding to a special row have outdegree 1, and all vertices corresponding to a top row have outdegree 0. In each row of M, only the rightmost 1 can be the end vertex of an edge in G, otherwise a Γ would be formed; and for this unique possible end vertex v, there is at most one edge in G which ends at v, the edge that starts from highest 1 below v in its column (if such a 1 exists). This implies that every vertex of H has indegree at most 1. These altogether mean that H consists of vertex-disjoint directed paths. The end vertices of the path components correspond to a top row, while the other vertices of H correspond to special rows. Consequently, the number of components of H is equal to the number of top rows in M, this number is denoted by m. Fig. 5 illustrates how H is obtained from G: H is the graph on the right; its red vertices correspond to top rows, and the black vertices correspond to special rows. (The labels of H will be discussed in the next paragraph, and the graph "projected G_M" was added to the figure for further purposes.)

The collection of pairs (R, C). We assign a pair (R, C) to each component P of H where R is the set of vertices (row indices) of P, and C is the set of column indices of top-1's in the top row corresponding to the endpoint of P. (On Fig. 5, the pairs (R, C) are written on the end vertices of the corresponding path components.) After doing this for all components, we get a collection of pairs $(R_1, C_1), \ldots, (R_m, C_m)$; we are left to define how to permute them to obtain a sequence. The obtained sequence will be an (n, k)-Callan sequence, because it has the required properties: The R_i's are obviously

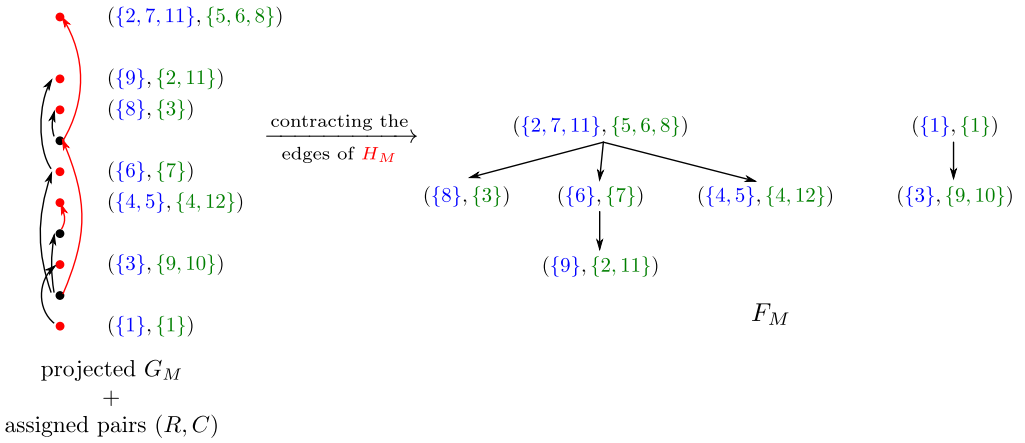

Fig. 6. The construction of the increasing forest F_M.

pairwise disjoint nonempty subsets of $\{1, \dots, n\}$ by construction; the C_i's are nonempty subsets of $\{1, \dots, k\}$ because every top row contains at least one top-1, and the C_i's are pairwise disjoint because every column of M contains at most one top-1 (all-0 columns have no top-1's, the other columns have exactly one top-1's).

The auxiliary graph F_M. Now we "project the regular edges of G horizontally". To this end, we define a new directed graph $F = F_M$ as follows: The vertices of F are the pairs $(R_1, C_1), \dots, (R_m, C_m)$; and for each *regular* edge e of G, there is an edge $\widetilde{e}$ in F, such that if e starts from row s and ends in row t, then $\widetilde{e}$ is from (R_i, C_i) to (R_j, C_j) where i and j are the unique indices for which s is contained in R_i and t is contained in R_j; and there are no other edges in F. It turns out that F is a rooted forest. We have seen above that the Γ-free property implies that for each row r of M, there is at most one edge in G that ends in r. This means that if e is a regular edge in G that starts from row s and ends in row t, then t is the smallest row index (lowest row) in the set R_j containing t, because there is no special edge ending in t. We also note that if s is contained in the set R_i, then $\min R_i \leqslant s < t = \min R_j$. So we can conclude that for each set R_j, there is at most one regular edge in G that ends in a row of R_j (this row can only be the lowest row, and there is at most one edge that ends in that row), implying that every vertex (R_i, C_i) in F has indegree at most 1. The previous discussion also showed that if $\widetilde{e}$ is an edge from (R_i, C_i) to (R_j, C_j) in F, then $\min R_i < \min R_j$. From these we can see that F is indeed a rooted forest (the edges are directed from parent to child), and moreover, F is an *increasing forest* on the set $\{(R_1, C_1), \dots, (R_m, C_m)\}$ equipped with the following total order $\prec$:

$$(R_i, C_i) \prec (R_j, C_j) \overset{\text{def}}{\iff} \min R_i < \min R_j.$$

(This order is total, because the R_i's are pairwise disjoint, so their smallest elements are pairwise distinct.) The construction of F is visualized in Fig. 6 (cf. Fig. 5).

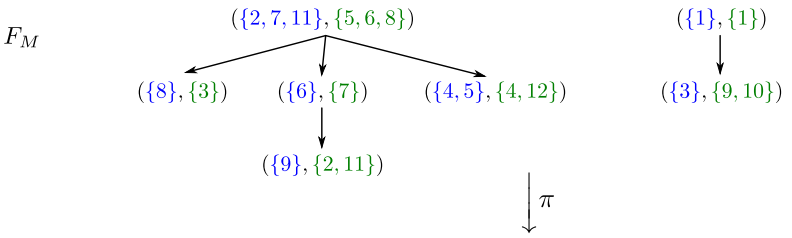

Fig. 7. Obtaining the (n,k)-Callan sequence from F_M.

The definition of ϕ. The bijection π of Lemma 6 converts the increasing forest F into a permutation $\pi(F)$ of the vertices $(R_1,C_1),\dots,(R_m,C_m)$. Finally, we set $\phi(M):=\pi(F)$; see also Fig. 7. We have already checked earlier that any permutation of the pairs $(R_1,C_1),\dots,(R_m,C_m)$ is an (n,k)-Callan sequence, so $\phi(M)\in\mathcal{S}$, i.e. the definition of ϕ is valid.

Some properties of ϕ. Before proving the bijectivity of ϕ, we summarize some properties of the construction:

(i) The rows in $\cup_{i=1}^m R_i$ are exactly those rows of M that contain at least one 1, the columns in $\cup_{i=1}^m C_i$ are exactly those columns of M that contain at least one 1.
(ii) In the row $\max R_i$ (from bottom), the top-1's are exactly in the columns of C_i, for $i=1,\dots,m$. There are no top-1's in other rows of M.
(iii) The edges of G_M are bijectively associated to the non-top 1's of M (an edge corresponds to the non-top 1 from which it starts).
(iv) If there is an edge e from row s to row t in G_M, then the start and end vertex of e are in the column of the rightmost 1 of the row t.
(v) The number of edges in H_M is equal to the number of special edges in G_M. (The "projection" $e\mapsto e'$ in the definition of H_M is a bijection between the two edge sets.)
(vi) The number of edges in F_M is equal to the number of regular edges in G_M. (The "projection" $e\mapsto\widetilde{e}$ in the definition of F_M is a bijection between the two edge sets.)
(vii) If (R_j,C_j) is a child of (R_i,C_i) in F_M, i.e. if there is an edge $\widetilde{e}$ from (R_i,C_i) to (R_j,C_j) in F_M, then the regular edge e of G_M that corresponds to $\widetilde{e}$ ends in row $\min R_j$ and starts in the highest row of R_i that is lower than $\min R_j$. (Such a row exists, because $\min R_i<\min R_j$ by the increasing property of F_M.)

The injectivity of the projection in (vi) was implicitly proved when we saw that every vertex in F_M has indegree at most 1. The last statement of (vii) follows from the facts that e cannot start from a higher row of R_i, because it is directed "upwards", and e cannot start from a lower row of R_i, because otherwise the regular edge e would be longer than the special edge starting from the same row.

Let $a = ((R_1^*, C_1^*), \dots, (R_l^*, C_l^*))$ be an arbitrary (n,k)-Callan sequence. We have to show that there is exactly one Γ-free $n \times k$ 0–1 matrix M^* for which $\phi(M^*) = a$.

The injectivity of ϕ. So we reconstruct M^* from a. The last step in constructing the image of a matrix was the application of the bijection π from Lemma 6. This step is invertible, so we can view a as $\pi^{-1}(a)$, an increasing forest on the set $\{(R_1^*, C_1^*), \dots, (R_l^*, C_l^*)\}$ with the total order $\prec$ defined above. Set $F^* := \pi^{-1}(a)$, and let H^* be the directed graph that is the vertex-disjoint union of directed paths P_i whose vertices are the elements of R_i^* in increasing order along the path $(i = 1, \dots, l)$. The point is that the top-1's of M^* are uniquely determined by property (ii); and we can uniquely determine the projection of the edges assigned to the non-top 1's, using property (v) for special edges and using properties (vi)–(vii) for regular edges. On a *projected edge* we mean a pair (s,t) of rows, which can be thought as "it starts from row s and ends in row t" $(s < t)$. Further, we can figure out the original column of the projected edge (s,t) once the row t is completely filled, by property (iv). A row r is completely filled if we know the positions of non-top 1's in that row (the top-1's are already determined); and for that, we have to project back the projected edges starting from r. For this reason, we may need first to fill the last rows of the children of (R_i^*, C_i^*) in F^*, where $R_i^* \ni r$. These lead to the following algorithm for constructing M^*:

- Start with an empty $n \times k$ matrix, then fill the rows of $\{1, \dots, n\} \setminus \cup_{i=1}^{l} R_i^*$ with 0's.
- Order the vertices of F^* in such a way that every non-root vertex precedes its parent (this can be done by post-order transversal, for example).
- Process the vertices of F^* in this order, and for vertex (R_i^*, C_i^*), fill the rows of R_i^* as follows $(i \in \{1, \dots, l\})$:
  (a) In row $\max R_i^*$, place 1's into the columns of C_i^*.
  (b) For each child (R_j^*, C_j^*) of (R_i^*, C_i^*): Let c denote the column of the rightmost 1 in row $\min R_j^*$ (that row is already filled), let r denote the highest row of R_i^* that is lower than $\min R_j^*$ (cf. property (vii), and we note that such an r exists by the increasing property of F^*), and place a 1 into position (r,c).
  (c) Fill the empty positions in row $\max R_i^*$ with 0's.
  (d) Let r_h be the h^{th} highest row (index) in R_i^*. For $h = 2, 3, \dots, |R_i^*|$, in row r_h place a 1 below the rightmost 1 of row r_{h-1}, and then fill the empty positions of r_h with 0's.

After the previous discussion, it should be clear that the obtained matrix M^* is the unique candidate to be the inverse of a. (We note that the outcome M^* of the algorithm does not depend on the actual order chosen in second step. We place at least one 1 to each row of $\cup_{i=1}^{l} R_i^*$, so the "rightmost 1" always makes sense in the algorithm.)

The surjectivity of ϕ. So we are left to show that M^* is Γ-free and that $\phi(M^*) = a$. In our task, the following observation will be helpful:

(viii) Let (R_i^*, C_i^*) be an arbitrary vertex of F^*, and let $J \subseteq \{1, \dots, l\}$ denote the set of indices j for which (R_j^*, C_j^*) is a vertex of $F^*[(R_i^*, C_i^*)]$ (cf. Notation 5). In M^* every 1 of a row of R_i^* is contained in a column of $\cup_{j\in J} C_j^*$. The rightmost 1 of row $\min R_i^*$ is in the rightmost column of $\cup_{j\in J} C_j^*$.

This can be proved by a routine induction on the height of (R_i^*, C_i^*) in F^*.

It is crucial to see that in step (b) when we place a 1 into position (r, c) we project back the edge $(R_i^*, C_i^*) \to (R_j^*, C_j^*)$ of F^* to a regular edge, while in step (d) we project back the edges of H^* to special edges. In order to see that this is what really happens in the algorithm, we have to check that there are no other 1's between the positions (r, c) and $(\min R_j^*, c)$ (the newly written 1 and the "rightmost 1") in the final M^*, otherwise the edge starting from (r, c) would not end at $(\min R_j^*, c)$, resulting $\phi(M^*) \neq a$. The analogous check for step (d) is also necessary. So we take a closer look at the way the 1's are placed in steps (a)–(d). When step (a) is applied, the newly written 1's are the first 1's appearing in their column. This is because, by property (viii), only those rows can contain 1 in a column of C_i^* that belong to an ancestor of (R_i^*, C_i^*), but the ancestors of (R_i^*, C_i^*) have not been considered yet. Then, applying property (viii) to the children of (R_i^*, C_i^*), we can see that the newly written 1's in step (b) has pairwise distinct columns, and different from the columns of C_i^*, too. This also means that we place $|R_i^*| - 1$ *new* 1's in step (d). For steps (b) and (d), the key observation is the following:

(∗) *The "rightmost* 1*" in the description of steps (b) and (d) is the lowest* 1 *in its column at the moment of placing the new* 1 *below it, for all possible indices* i, j*.*

Before proving this, we note that (∗) and the observation made on step (a) imply that *the algorithm fills the* 1*'s of each column from top to bottom*. Now we prove (∗) by induction on the progress of the algorithm. Suppose, by contradiction, that (∗) does not hold for step (b) for some i, j; and let u denote the "rightmost 1" in question, and let v denote the "newly written 1". In other words, suppose that there is a 1 below u, denoted by w, when we want to write v. We can assume that there are no other 1's between u and w. As seen before, the substeps of step (b) involve pairwise distinct columns c for the children of this given (R_i^*, C_i^*), and $c \notin C_i^*$, so v is the first 1 in the column of u that is written into a row of R_i^*. Since u is in row $\min R_j^*$, w is not in R_j^* (and not in R_i^*, neither); let R_k^* be the set containing w ($j \neq k \neq i$). By the induction hypothesis and an easy analysis of the algorithm, we know that w was placed below u in a former step (b), which means that (R_j^*, C_j^*) has two parents in F^*, (R_i^*, C_i^*) and (R_k^*, C_k^*), contradiction. After processing the steps (a)–(c), all the 1's written in a row of R_i^* are the lowest ones in their column, as just seen. From this, it is very easy to conclude that (∗) holds for step (d) as well.

We prove now that M^* is Γ-free, i.e. $M^* \in \mathcal{M}$. Suppose, by contradiction, that M^* contains three 1's in a Γ-configuration, and let u, v and w denote the upper-left, lower-left and the upper-right one, respectively: $\begin{smallmatrix} u & w \\ v & \end{smallmatrix}$. We can assume that v is chosen so that there are no 1's between u and v in M^*. We saw in the previous paragraph that the algorithm fills the 1's of each column from top to bottom in the way described in (∗),

which means that v was placed below u in a step (b) or (d). But this is a contradiction, because u is not the rightmost 1 of its row (and the whole row of u has been already filled, when v is placed).

Now we sketch why $\phi(M^*) = a$. As we underlined earlier, for steps (b) and (d), there are no other 1's between the "rightmost 1" and the "newly written 1" in the final M^*. This is ensured by $(*)$ at the moment after writing the new 1, and it will not change later neither, because the algorithm fills the 1's from top to bottom per column. From this it is easy to see that in step (b) we indeed project back the edges of F^* to regular edges, more precisely, we ensure that the regular edges (of G_{M^*}) starting from a row of R_i^* become, in F_{M^*}, the edges of F^* starting from (R_i^*, C_i^*). (We leave the reader to check that the newly placed 1's in step (b) are start vertices of *regular* edges in the final M^*.) Similarly, in step (d) we indeed project back the edges of H^* to special edges. (We leave the reader to check that the newly placed 1's in step (d) are start vertices of *special* edges in the final M^*.) Finally, we note that since the 1's are filled from top to bottom per column, hence the 1's placed in step (a) will be top-1's in the final M^*, too – which is also required for M^* being the inverse of a. Some minor details were skipped, but it is easy to see now that $H_{M^*} = H^*$, the pairs assigned to M^* are $\{(R_i^*, C_i^*)\}_{i=1}^{l}$, and $F_{M^*} = F^*$, which lead to $\phi(M^*) = a$, as stated. □

Recall from Section 1 that non-ambiguous forests correspond to those Γ-free 0–1 matrices that have no all-0 rows and columns. So in order to count non-ambiguous forests, one needs to enumerate these Γ-free matrices. This was done in [5], and it is also an easy corollary of our previous construction.

Corollary 8. *The number of non-ambiguous forests with characteristic matrix of size $n \times k$ is*

$$\sum_{m=0}^{\min(n,k)} (m!)^2 \left\{ {n \atop m} \right\} \left\{ {k \atop m} \right\}.$$

Proof. By the definition of non-ambiguous forests, we have to count the Γ-free $n \times k$ 0–1 matrices without all-0 rows and columns. Let $\mathcal{N}$ denote the set of these matrices. In the proof of Theorem 2, ϕ establishes a bijective correspondence between Γ-free $n \times k$ 0–1 matrices and (n, k)-Callan sequences. Property (i) of ϕ shows that the restriction of ϕ to $\mathcal{N}$ is a bijection between $\mathcal{N}$ and the set of those (n, k)-Callan sequences $(R_1, C_1), \ldots, (R_m, C_m)$ for which $\{R_1, \ldots, R_m\}$ is a partition of $\{1, \ldots, n\}$ and $\{C_1, \ldots, C_m\}$ is a partition of $\{1, \ldots, k\}$. These sequences are clearly counted by the sum in the statement. □

3. The generating function of Γ-free matrices

We can use the bijection to derive easily the generating function of Γ-free matrices. We introduce some notations. Let $\mathcal{G}$ denote the set of Γ-free matrices. For $m \in \mathcal{G}$ let

$r(m)$ denote the number of rows, $c(m)$ the number of columns, $r_e(m)$ the number of empty rows, $c_e(m)$ the number of empty columns, and $r_t(m)$ denote the number of top rows, rows that contain top 1's. Let $M(x, y, a, b, t)$ be the generating function of Γ-free matrices defined as follows:

$$M(x, y, a, b, t) = \sum_{m \in \mathcal{G}} t^{r_t(m)} a^{r_e(m)} b^{c_e(m)} \frac{x^{r(m)}}{r(m)!} \frac{y^{c(m)}}{c(m)!}.$$

We have the next theorem.

Corollary 9. *The generating function of* Γ*-free matrices is*

$$M(x, y, a, b, t) = \frac{e^{ax} e^{by}}{1 - t(e^x - 1)(e^y - 1)}.$$

Proof. We have seen that there is a bijection between Γ-free matrices of size $n \times k$ and (n, k)-Callan sequences. Moreover, we have seen that $\cup_{i=1}^{m} R_i$ (resp. $\cup_{i=1}^{m} C_i$) corresponds to the set of non-empty rows (resp. columns), and each top row corresponds to a pair (R_i, C_i). We can construct an (n, k)-Callan sequence as a sequence (permutation) of pairs of non-empty sets from $\{1, \dots, n\}$ and $\{1, \dots, k\}$ with two further sets: the set of remaining elements from $\{1, \dots, n\}$ that may be empty; and the set of remaining elements from $\{1, \dots, k\}$ that may also be empty. We follow the standard technique of symbolic method and use the usual notations of [8] for combinatorial constructions without explicitly defining these here. The construction for Callan sequences described above is symbolically formulated as follows:

$$(\mathrm{SEQ}(\gamma \mathrm{SET}_{>0}(\mathcal{X}) \times \mathrm{SET}_{>0}(\mathcal{Y}))) \times (\mathrm{SET}(\alpha \mathcal{X})) \times (\mathrm{SET}(\beta \mathcal{Y})),$$

where $\mathcal{X}$ is the atomic class from $\{1, \dots, n\}$, $\mathcal{Y}$ is the atomic class from $\{1, \dots, k\}$, α marks the empty rows, β marks the empty columns, and γ marks the top rows. Based on the correspondence: $\mathcal{X} \to x$, $\mathcal{Y} \to y$, $\alpha \to a$, $\beta \to b$, and $\gamma \to t$ the construction translates to the formula given in the theorem. □

4. Proof of Theorem 4

This section is devoted to the proof of Theorem 4. Our construction is based on the fact that pairs of permutations of $\{1, \dots, n\}$ with no common rise can be encoded with certain labeled rooted forests. We present this as a lemma, but we think it is interesting in its own right.

The set of permutations of $\{1, \dots, n\}$ is denoted by S_n; and for a permutation $\eta = (\eta_1, \dots, \eta_n) \in S_n$, we denote by P_η the set $\{(1, \eta_1), \dots, (n, \eta_n)\}$.

Recall Convention 3 and Fig. 1 where the roots are colored red: The children of an "element 1" (vertex u) in a non-ambiguous forest are the lowest 1 above u (if such

a 1 exists) and the rightmost 1 on the left side of u (if such a 1 exists). We often mix the (characteristic) matrix terminology and the graph terminology for non-ambiguous forests as in the previous sentence, if it does not lead to confusion.

Let $M = \chi_A$ be the characteristic matrix of a non-ambiguous forest A. Analogously to the notion of top-1's introduced in Section 2, we say that an element 1 in M is a *leading*-1, if it is the leftmost 1 in its row. Clearly, an element 1 in M has exactly one child in A if and only if the element is either a top-1 but not a leading-1 or it is a leading-1 but not a top-1. So A is complete if and only if there is no such element, i.e. the set of top-1's coincides with the set of leading-1's in M. Since M has no all-0 rows and columns, there is exactly one leading-1 in each row and there is exactly one top-1 in each column. So the top-1's can coincide with the leading-1's only if M is a square matrix. Moreover, if A is complete, then the leaves are exactly the top-1's / leading-1's, because a leaf is always a top-1, and now each top-1 is a leaf because it is also a leading-1. This means that in case of complete A, every row and every column contains exactly one leaf (the leading-1 or top-1), and that the number of rows / columns is equal to the number of leaves. We have just proved the following.

Lemma 10. *The characteristic matrices of* complete *non-ambiguous forests with n leaves are exactly those Γ-free* 0–1 *matrices without all-*0 *rows and columns that have size $n \times n$ and in which every top-*1 *is a leading-*1*. In such a matrix M the leaves are exactly the top-*1*'s, and the set of positions of leaves/top-*1*'s is P_η for some $\eta \in S_n$.*

A pair $(\alpha, \beta) \in S_n \times S_n$ will always be identified with the sequence $(a_1, b_1), \dots, (a_n, b_n)$, where $\alpha = (a_1, \dots, a_n)$ and $\beta = (b_1, \dots, b_n)$. If we want to emphasize this, we write $(\alpha, \beta)^T$ for denoting the sequence. In this setting, the pair (α, β) has *common rise* if there are two consecutive elements of $(\alpha, \beta)^T$ for which $(a_i, b_i) < (a_{i+1}, b_{i+1})$, where '$<$' stands for "less in both coordinates". For an arbitrary $(\alpha, \beta) \in S_n \times S_n$, the set of elements of $(\alpha, \beta)^T$ is clearly equal to a set P_η for some $\eta \in S_n$.

Now fix an n. We group the complete non-ambiguous forests with n leaves by the set of their leaves (i.e. the set of positions (i, j) of the leaves) and group the members of $S_n \times S_n$ without common rise by the set of their elements – considering the members of $S_n \times S_n$ as n-element sequences. From the discussion above, each of these groups can be described by a permutation of S_n. Hence, Theorem 4 clearly follows from the following theorem:

Theorem 11. *Fix an arbitrary $\eta \in S_n$. Let $\mathcal{C}$ be the set of those complete non-ambiguous forests (with n leaves) whose set of leaves is P_η. Let $\mathcal{P}$ be the set of those pairs $(\alpha, \beta) \in S_n \times S_n$ that have no common rise and for which the set of elements of $(\alpha, \beta)^T$ is P_η, i.e. $\mathcal{P}$ is the set of permutations of P_η with no common rise (in the sense above). Then there exists a bijection between $\mathcal{C}$ and $\mathcal{P}$, thus $|\mathcal{C}| = |\mathcal{P}|$ (Fig. 8).*

Proof. Let $\eta = (\eta_1, \dots, \eta_n)$. We think the elements of $\mathcal{C}$ as $n \times n$ 0–1 matrices, as characterized in Lemma 10. Recall the construction of the proof of Theorem 2.

$\mathcal{C}$: $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$ $\begin{matrix} 3 \\ 2 \\ 1 \end{matrix}$ $\begin{matrix} 3 & 2 & 1 \end{matrix}$ $\quad$ $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}$ $\begin{matrix} 3 \\ 2 \\ 1 \end{matrix}$ $\begin{matrix} 3 & 2 & 1 \end{matrix}$ $\quad$ $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 1 \end{pmatrix}$ $\begin{matrix} 3 \\ 2 \\ 1 \end{matrix}$ $\begin{matrix} 3 & 2 & 1 \end{matrix}$ $\quad$ $\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}$ $\begin{matrix} 3 \\ 2 \\ 1 \end{matrix}$ $\begin{matrix} 3 & 2 & 1 \end{matrix}$

$\mathcal{P}$: $\big((3,2),(2,1),(1,3)\big)$ $\big((2,1),(1,3),(3,2)\big)$ $\big((3,2),(1,3),(2,1)\big)$ $\big((1,3),(3,2),(2,1)\big)$

Fig. 8. The sets $\mathcal{C}$ and $\mathcal{P}$ for $n = 3$, $\eta = (3, 1, 2)$.

As a first step, we apply the injective map $M \mapsto F_M = \pi^{-1}(\phi(M))$ to the matrices $M \in \mathcal{C}$, where F_M, π, and ϕ are defined in Section 2. This map establishes a bijective correspondence between the matrices of $\mathcal{C}$ and the increasing forests in the set $\widetilde{\mathcal{C}} := \{F_M : M \in \mathcal{C}\}$. Now we are going to give a simple description of $\widetilde{\mathcal{C}}$.

By Lemma 10, we know that in every matrix of $\mathcal{C}$, the set of top-1's is P_η. Let $\mathcal{D}$ denote the set of Γ-free $n \times n$ 0–1 matrices in which the set of top-1's is P_η. Clearly, $\mathcal{D} \supseteq \mathcal{C}$, but not every matrix of $\mathcal{D}$ corresponds to a *complete* non-ambiguous forest. Now we pick an arbitrary matrix $M \in \mathcal{D}$, and examine the increasing forest F_M. We know that each row of M is a top row: the row i has exactly one top-1, which is in column η_i, by the definition of $\mathcal{D}$. This means that the vertex set of F_M is P_η. (More precisely, the vertices of F_M have the form $(\{i\}, \{\eta_i\})$, but we can leave the braces.) As there are no special edges in G_M (see Section 2), every non-top 1 of M contributes to the edges of F_M. In this special case, F_M can be obtained from M as follows: View M as a non-ambiguous forest (with vertices and edges), project it horizontally, then label the vertex corresponding to the projected row i with (i, η_i), for $i = 1, \ldots, n$, and orient the edges upwards. It should be clear from the proof of Theorem 2 that the map $M \mapsto F_M$ is a bijection between $\mathcal{D}$ and the set of increasing forests on vertex set P_η. We denote the latter set by $\widetilde{\mathcal{D}}$. The total order $\prec$ on vertex set P_η is simply the order by first coordinate, i.e. now the condition "increasing" means that every child must have greater first coordinate than its parent has.

We know that $\widetilde{\mathcal{C}} \subseteq \widetilde{\mathcal{D}}$. An increasing forest $F^* \in \widetilde{\mathcal{D}}$ is in $\widetilde{\mathcal{C}}$ if and only if the corresponding matrix $M^* \in \mathcal{D}$ (for which $F_{M^*} = F^*$) is in $\mathcal{C}$. By Lemma 10, the matrix M^* of $\mathcal{D}$ is in $\mathcal{C}$ if and only if every top-1 is a leading-1 in M^*, in other words, iff in every row of M^* the non-top 1's are on the right side of the top-1 in that row. The column index of a non-top 1 can be read off from F^*. In the proof of Theorem 2, step (b) of the inverse construction describes how the non-top 1's are placed to M^* (step (d) does not place any 1's for $F^* \in \widetilde{\mathcal{D}}$): For an arbitrary row r, the non-top 1's are associated to the children of (r, η_r) in F^*. Namely, for each child (s, η_s) of (r, η_r), a 1 is placed to row r into the column of the rightmost 1 of row s, and that column is the minimum of the second coordinates of the vertices in the subtree $F^*[(s, \eta_s)]$, applying the last statement of property (viii) to the child (s, η_s). (Cf. Notation 5 and Convention 3.) We conclude that the increasing forest $F^* \in \widetilde{\mathcal{D}}$ is in $\widetilde{\mathcal{C}}$ if and only if for every vertex (a, b) of F^* and for every child (c, d) of (a, b), there exists a vertex of $F^*[(c, d)]$ with second coordinate less than b.

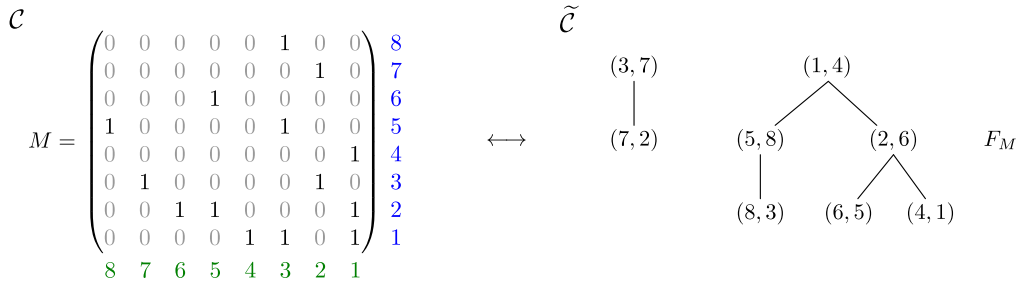

Fig. 9. A complete non-ambiguous forest and the corresponding properly labeled forest.

We introduce a term for describing the forests of $\widetilde{\mathcal{C}}$. We say that F is a *properly labeled forest* on vertex set P_η (for some $\eta \in S_n$), if F is an (unordered) rooted forest on vertex set P_η, satisfying that whenever (a, b) is the parent of (c, d) in F, then $a < c$ and there exists a vertex $(x, y) \in F[(c, d)]$ such that $b > y$. (See Fig. 9 for an example.) We can summarize the above investigations as $\widetilde{\mathcal{C}}$ is the set of properly labeled forests on P_η. So it is enough to prove the following lemma. □

Lemma 12. *For a given $\eta \in S_n$, let $\widetilde{\mathcal{C}}$ denote the properly labeled forests on vertex set P_η, and let $\mathcal{P}$ denote the set of permutations of P_η with no common rise. Then there exists a bijection between $\widetilde{\mathcal{C}}$ and $\mathcal{P}$, thus $|\widetilde{\mathcal{C}}| = |\mathcal{P}|$.*

Proof. We begin with some conventions. For $(a, b), (c, d) \in P_\eta$, let

$$(a, b) <_1 (c, d) \overset{\text{def}}{\iff} a < c;$$

$$(a, b) <_2 (c, d) \overset{\text{def}}{\iff} b < d.$$

In this proof we will work with increasing forests on vertex set $(P_\eta, <_1)$. We follow the conventions introduced at the beginning of Section 2 (cf. Fig. 2): We always list the children of a given parent of a forest F in decreasing order from left to right with respect to the order $<_1$, and the phrase "*leftmost/first* child" refers to this list (i.e. it means the child of a given parent with biggest first coordinate). The tree components of F are also listed in the decreasing order of their roots with respect to $<_1$.

As a first step, we apply π^{-1} to the permutations of $(P_\eta, <_1)$, where π is the bijection introduced in the proof of Lemma 6. (So we consider only the first coordinates with the natural order when building the forest structure from a permutation, see also Fig. 10.) By the lemma, π^{-1} establishes a bijective correspondence between the permutations of $(P_\eta, <_1)$ and the increasing forests on $(P_\eta, <_1)$. It is known, or an easy analysis of π shows, that for any permutation (sequence) $S = (\mathbf{s}_1, \dots, \mathbf{s}_n)$ of P_η and any $i \in \{1, \dots, n-1\}$, the inequality $\mathbf{s}_i <_1 \mathbf{s}_{i+1}$ holds if and only if $\mathbf{s}_{i+1}$ is the *leftmost* child of $\mathbf{s}_i$ in the increasing forest $\pi^{-1}(S)$. This implies that S has no common rise if and only if, for every vertex u

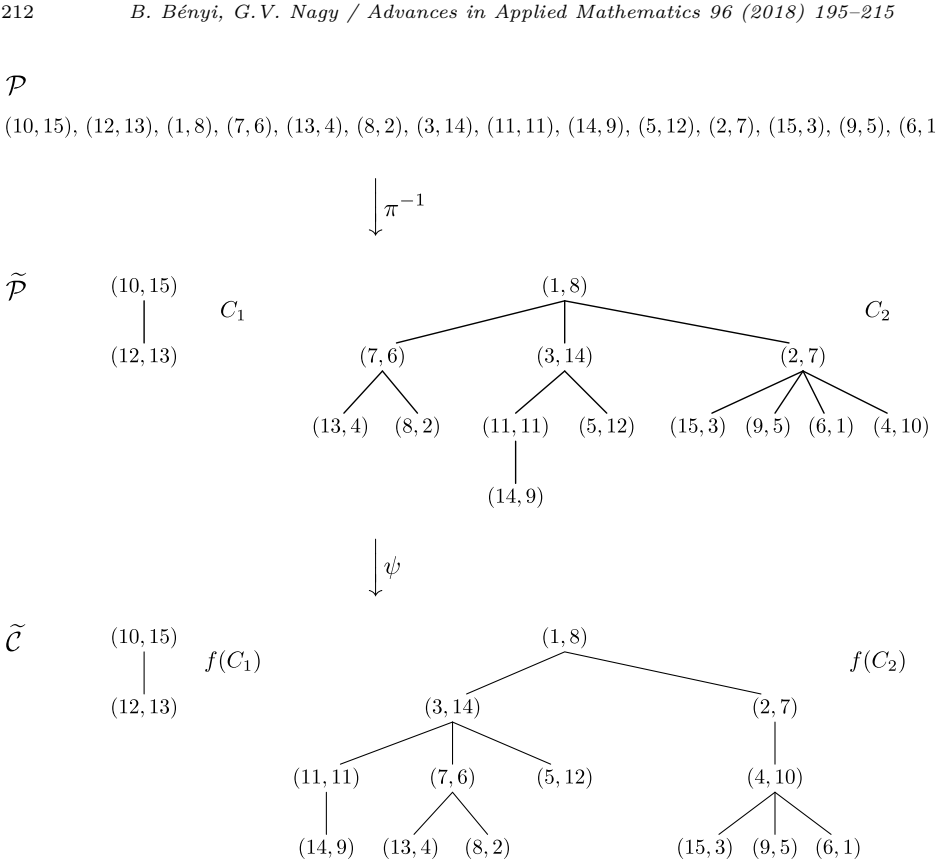

Fig. 10. Illustration of $\pi^{-1}|_{\mathcal{P}}$ and ψ (for $n = 15$).

of $\pi^{-1}(S)$, the leftmost child of u (if u is a non-leaf) has smaller second coordinate than u has. So $\widetilde{\mathcal{P}} := \{\pi^{-1}(S) : S \in \mathcal{P}\}$ is the set of those (unordered) rooted forests F on vertex set P_η which satisfy the following conditions:

(c1) whenever vertex v is a child of vertex u in F, then $u <_1 v$, i.e. F is increasing with respect to $<_1$;
(c2) and whenever vertex v is the *leftmost* child of vertex u in F, then $v <_2 u$.

We say that a forest F is *leftmost-valid*, if it satisfies conditions (c1)–(c2). So $\widetilde{\mathcal{P}}$ is the set of leftmost-valid forests on P_η.

The definition of the main bijection. In order to complete the proof, we give a bijection $\psi\colon \widetilde{\mathcal{P}} \to \widetilde{\mathcal{C}}$. Both the notion of leftmost-valid forest and the notion of properly labeled forest can be extended for any vertex set $V \subseteq P_\eta$, without any modification. First we define a bijective conversion function f from the set of leftmost-valid *trees* (one-component forests) to the set of properly labeled *trees*, such that $f(T)$ has the same vertex set and root as T, for any leftmost-valid tree T. Then we can define ψ. For an arbitrary $F \in \widetilde{\mathcal{P}}$, if F has (the leftmost-valid tree) components $C_1, \ldots, C_m$, then $\psi(F)$ is defined to be

the vertex-disjoint union of the properly labeled trees $f(C_1), \dots, f(C_m)$. As f keeps the vertex set, $\psi(F) \in \widetilde{\mathcal{C}}$.

Now we give the (recursive) definition of f. We note that for any leftmost-valid (resp. properly labeled) tree T on V, $T[v]$ is clearly a leftmost-valid (resp. properly labeled) tree for any $v \in V$, cf. Notation 5. It is recommended to follow the conversion of C_2 on Fig. 10. For a leftmost-valid tree T on vertex set $V \subseteq P_\eta$, we define $f(T)$ as follows.

- If T has one vertex, then $f(T) := T$.
- Otherwise, let r be the root of T, and let $v_1, \dots, v_k$ be the children of r in $<_1$-decreasing order. Set $T_i := f(T[v_i])$ for $i = 1, \dots, k$, and consider the *sequence*

$$T_1,\, T_2,\, \dots,\, T_k. \tag{2}$$

  Find the smallest index $i \in \{1, \dots, k\}$, if such an i exists, for which no vertex of T_i has smaller second coordinate than r has. (We say that T_i is the leftmost *bad* tree.) We note that $i \neq 1$, as it will be justified later. Then remove the elements $T_1, \dots, T_i$ from (2), and add a new first element T_1', where T_1' is the rooted tree obtained from $T_1, \dots, T_i$ by joining the roots of $T_1, \dots, T_{i-1}$ (as new children) to the root of T_i (as parent/root). In this way we obtain a new sequence

$$T_1',\, T_2',\, \dots,\, T_{k-i+1}'; \tag{3}$$

  where $T_j' = T_{j+i-1}$ for $j \geqslant 2$. We call this process *merging*. Then do the same for (3): find the leftmost bad tree, merge. Then repeat this for the new sequence, and so on, stop when no such index i (bad tree) was found. We note that the process terminates, because the length of the sequence strictly decreases in each merging step ($i \neq 1$). We end up with a sequence $\widetilde{T}_1, \dots, \widetilde{T}_l$. Finally, $f(T)$ is defined to be the tree with root r that is obtained by joining the vertex r (as new root) to the roots of $\widetilde{T}_1, \dots, \widetilde{T}_l$.

Now we justify why $i \neq 1$, i.e. why the first tree in the actual tree sequence always has a vertex which is smaller than r in the second coordinate. This is true for the initial sequence (2), because v_1, the root of T_1, has smaller second coordinate than r has, by the leftmost-validity of T and the fact that f keeps the roots. And this vertex v_1 will keep staying in the first tree of the sequence after the mergings, too.

The fact that $f(T)$ has the same vertex set and root as T can be verified by an easy induction.

Now we show why $f(T)$ is properly labeled. By induction, every tree of the initial sequence (2) is properly labeled. It is straightforward to see that this property is kept after each merging. We only have to check the new parent–child connections between the root of the bad tree and its new children, the roots of the preceding trees. The monotonicity condition on the first coordinates is satisfied because the initial roots $v_1, \dots, v_k$ are in

$<_1$-decreasing order. The condition on the second coordinates is satisfied because if the root of the bad tree is v_b, then $r <_2 v_b$ (as v_b is vertex of a bad tree), while in every preceding tree there exists a vertex u with $u <_2 r$ (as they are good trees), so a vertex u with $u <_2 v_b$, as needed. In the final stage we have the properly labeled good trees $\widetilde{T}_1, \dots, \widetilde{T}_l$, from which it is pretty obvious that $f(T)$ is properly labeled (after checking the root r).

The inverse of f. Now we sketch why ψ is a bijection. It is clear that it is enough to show that for any fixed vertex set $V \subseteq P_\eta$, the function f (or more precisely, its restriction) is a bijection between the set of leftmost-valid trees on V and the set of properly labeled trees on V. This can be done by induction on the size of V. Pick an arbitrary properly labeled tree T^* on vertex set V, where $|V| \geqslant 2$. Let r be the root of T^*. The children of a given parent are listed in $<_1$-decreasing order (the pre-order transversal follows this order in the next step).

- Find the first vertex v_1 in the pre-order transversal of T^* (see the proof of Lemma 6) for which $v_1 <_2 r$. As T^* is properly labeled, such a v_1 exists in $T^*[u]$, where u is the leftmost child of r.
- We define a list (sequence) $\mathfrak{L}$ of subtrees. As an initial step, we define the first element of $\mathfrak{L}$ to be $T^*[v_1]$.
- Let $w_1, \dots, w_m$ be those siblings of v_1 in T^*, listed in $<_1$-decreasing order, for which $T^*[w_i]$ has a vertex with smaller second coordinate than the second coordinate of r. We refer to these siblings as *good* siblings. (The good siblings are on the right side of v_1, because v_1 was found by pre-order transversal.) Add the trees $T^*[w_1], \dots, T^*[w_n]$ to the end of $\mathfrak{L}$ in this order. Then let v_2 be the parent of v_1 in T^*, and let $T^-_{v_2}$ denote the tree obtained from $T^*[v_2]$ by deleting the subtrees $T^*[v_1]$, $T^*[w_1], \dots, T^*[w_n]$ from it. Add $T^-_{v_2}$ to the end of $\mathfrak{L}$. Repeat this step for v_2 instead of v_1, and then for the parent v_3 of v_2, and so on, until we reach to the point when $v_i = u$. At that point, all siblings of $v_i = u$ are good, due to the fact that T^* is properly labeled. Add the subtrees $T^*[w]$ to $\mathfrak{L}$ for the siblings w of u from left to right as above, which finishes this step.
- We end up with a list $\mathfrak{L}$ of trees: $L_1, \dots, L_t$. Finally, $\widehat{T}$ is defined to be the tree with root r that is obtained by joining the vertex r (as new root) to the roots of $f^{-1}(L_1), \dots, f^{-1}(L_t)$, where the unique inverse images $f^{-1}(L_i)$ come from the induction hypothesis.

We claim that $\widehat{T}$ is the unique leftmost-valid tree on vertex set V for which $f(\widehat{T}) = T^*$, proving the bijectivity of f. The details are left to the reader. □

Second author's acknowledgments

This research was supported by the EU-funded Hungarian grant EFOP-3.6.1-16-2016-00008.

References

[1] N. Aisbett, On the poset of vector partitions, arXiv:1505.01996v2.
[2] J.C. Aval, A. Boussicault, M. Bouval, M. Silimbani, Combinatorics of non-ambiguous trees, Adv. in Appl. Math. 56 (2014) 78–108.
[3] J.C. Aval, A. Boussicault, B. Delcroix-Oger, F. Hivert, P. Laborde-Zubieta, Non-ambiguous trees: new results and generalization, arXiv:1511.09455v1.
[4] B. Bényi, P. Hajnal, Combinatorics of poly-Bernoulli numbers, Studia Sci. Math. Hungar. 52 (4) (2015) 537–558.
[5] B. Bényi, P. Hajnal, Combinatorial properties of poly-Bernoulli relatives, Integers 17 (2017) A31.
[6] C.R. Brewbaker, A combinatorial interpretation of the poly-Bernoulli numbers and two Fermat analogues, Integers 8 (2008) A02.
[7] L. Carlitz, R. Scoville, T. Vaughan, Enumeration of pairs of permutations and sequences, Bull. Amer. Math. Soc. 80 (5) (1974) 881–884.
[8] P. Flajolet, R. Sedgewick, Analytic Combinatorics, Cambridge University Press, Cambridge, 2009.
[9] Z. Füredi, P. Hajnal, Davenport Schinzel theory of matrices, Discrete Math. 103 (1992) 233–251.
[10] E.Y. Jin, Heaps and two exponential structures, European J. Combin. 54 (2016) 87–102.
[11] H.K. Ju, S. Seo, Enumeration of $(0,1)$-matrices avoiding some 2×2 matrices, Discrete Math. 312 (16) (2012) 2473–2481.
[12] S. Kitaev, T. Mansour, A. Vella, Pattern avoidance in matrices, J. Integer Seq. 8 (2005) A05.2.2.
[13] N.J.A. Sloane, The on-line encyclopedia of integer sequences, http://oeis.org.
[14] R. Stanley, Enumerative Combinatorics, vol. I, 2nd edition, Cambridge University Press, Cambridge, 2012.