A representation on the labeled rooted forests

Mahir Bilen Can

August 14, 2017

Abstract

We consider the conjugation action of symmetric group on the semigroup of all partial functions and develop a machinery to investigate character formulas and multiplicities. By interpreting these objects in terms of labeled rooted forests, we give a characterization of the labeled rooted trees whose S_n orbit afford the sign representation. Applications to rook theory are offered.

Keywords: Nilpotent partial transformations, labeled rooted trees, symmetric group, plethysm.
MSC: 05E10, 20C30, 16W22

1 Introduction

Our goal in this paper is to contribute to the general field of combinatorial representation theory by using some ideas from semigroup theory, rook theory, as well as graph theory. Classical rook theory is concerned with the enumerative properties of file and rook numbers in relation with other objects of mathematics [3]. In particular, enumeration of functions satisfying various constraints falls into the scope of rook theory. Here, we focus on partial functions, also known as partial transformations, on $[n] := \{1, \dots, n\}$ with the property that the associated graph of the function is a labeled rooted forest.

There is an obvious associative product on the set of all partial transformations on $[n]$; the composition $f \circ g$ of two partial transformations f and g is defined when the domain of f intersects the range of g. The underlying semigroup, denoted by $\mathcal{P}_n$ is called the partial transformation semigroup [6]. Of course, the identity map on $[n]$ is a partial transformation whence $\mathcal{P}_n$ is a monoid. Moreover, the subset $\mathcal{R}_n \subset \mathcal{P}_n$ consisting of injective partial transformations forms a submonoid of $\mathcal{P}_n$. $\mathcal{R}_n$ is known as (among combinatorialists) the rook monoid since its elements have interpretations as non-attacking rook placements on the "chessboard" $[n] \times [n]$ (see [8, 16]). It has a central place in the structure theory of reductive algebraic monoids. See [22, 23].

The symmetric group S_n is the group of invertible elements in both of the monoids $\mathcal{P}_n$ and $\mathcal{R}_n$. In this work, we compute the decompositions of certain representations of S_n on "nilpotent partial transformations." Since there is no obvious 0 element in $\mathcal{P}_n$, we explain this using a larger monoid. Let $\mathcal{F}ull_n$ denote the *full transformation semigroup* which consists of all maps from $\{0, 1, \dots, n\}$ into $\{0, 1, \dots, n\}$. The partial transformation semigroup $\mathcal{P}_n$ is canonically isomorphic to the subsemigroup $\mathcal{P}_n^* \subset \mathcal{F}ull_n$ consisting of elements $\alpha : \{0, 1, \dots, n\} \to \{0, 1, \dots, n\}$ such that $\alpha(0) = 0$. An element $g \in \mathcal{F}ull_n$ is called nilpotent if there exists a sufficiently large $k \in \mathbb{N}$ such that for any $x \in \{0, 1 \dots, n\}$, $g^k(x) = g(g(\cdots(x)\cdots)) = 0$. An element f of $\mathcal{P}$ is called nilpotent if, under the canonical identification $f \mapsto g = g_f$ of $\mathcal{P}$ with $\mathcal{P}^*$, the corresponding full transformation g is nilpotent.

The sets of nilpotent elements of $\mathcal{P}_n$ and $\mathcal{R}_n$ are denoted by $\mathrm{Nil}(\mathcal{P}_n)$ and $\mathrm{Nil}(\mathcal{R}_n)$, respectively. There is a beautiful way of representing, in terms of graphs, of the elements of these sets of nilpotent transformations. To build up to it, we first mention some useful alternative ways of representing the elements of $\mathcal{P}_n$.

Recall that a *partial transformation* is a function $f : A \to [n]$ that is defined on a subset A of $[n]$. We write the data of f as a sequence $f = [f_1, f_2, \dots, f_n]$, where $f_i = f(i)$ if $i \in A$, and $f_i = 0$ otherwise. Equivalently, f is given by the matrix $f = (f_{i,j})_{i,j=1}^n$ whose entries are defined by

$$f_{i,j} = \begin{cases} 1 & \text{if } f(i) = j; \\ 0 & \text{otherwise.} \end{cases} \tag{1.1}$$

The ordinary matrix multiplication of two such matrices corresponds to the compositional product of the corresponding partial transformations.

Finally, a more combinatorial way of representing $f \in \mathcal{P}_n$ is described as follows. Starting with n labeled vertices (labeled by the elements of $[n]$), if $f(i) = j$, then we connect the vertex with label i, by an outgoing directed edge, to the vertex with label j. The resulting graph is called the digraph of f. In Figure 1.1 we depict three different representations, including the digraph, of the partial transformation $f = [3, 3, 5, 0, 5, 1]$.

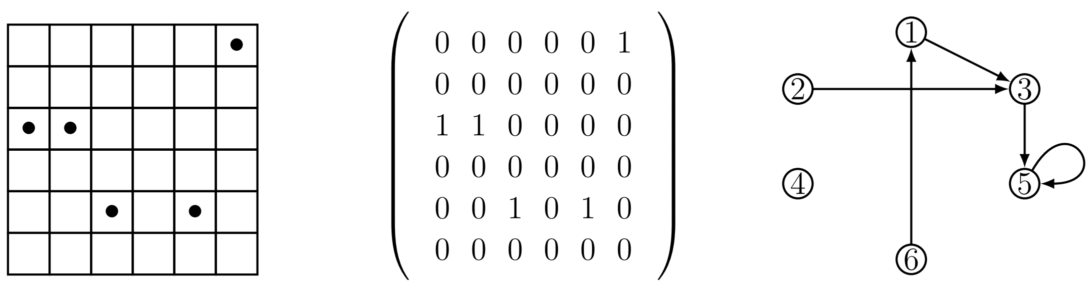

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Figure 1.1: Different presentations of the partial transformation $f = [3, 3, 5, 0, 5, 1]$.

The use of digraphs in rook theory goes back to Gessel's creative work [9]. This approach is taken much afar by Haglund in [12] and Butler in [2]. As far as we are aware of, the nilpotent rook placements in combinatorics made its first appearance in Stembridge and Stanley's influential work [25] on the immanants of Jacobi-Trudi matrices, where, essentially, the authors consider only those rook placements fitting into a staircase shape board. However, Stembridge and Stanley do not pursue the representation theoretic properties as we do here.

Following the terminology of [3], we call an element f of $\mathcal{P}_n$ a k-file placement if the number of nonzero entries of the matrix representation of f is k, and similarly, we call an element f of $\mathcal{R}_n$ a k-rook placement if the rank of the matrix representation of f is k. The basic observation that our paper builds on is that the symmetric group acts on nilpotent k-file placements as well as on the nilpotent k-rook placements. In fact, there is a more general statement from semigroup theory: the unit group of a monoid (with 0) acts on the set of nilpotent elements of the semigroup. Here we focus on the partial transformation monoid and the rook monoid. The main reason for confining ourselves only to these two special semigroups is twofold. First of all, the resulting objects from our investigations, namely, the labeled rooted forests, has significance not only in modern algebraic combinatorics ([7, 11, 13, 14], see [1], also) but also in the theory of classical transformation semigroups (see [6]). Secondly, much studied but still mysterious plethysm operation ([19, 15]) from representation theory has a very concrete combinatorial appearance in our work. Furthermore, rooted trees and forests play a very important role for computer science (see [17]). There are plenty of other reasons to focus on these objects for statistical and probabilistic purposes.

Now we are ready to give a brief overview of our paper and state our main results. In Section 2, we introduce the necessary notation and state some of the results that we use in the sequel. The purpose of Section 3 is to give a count of the number of file placements (partial transformations) according to the sizes of their domains. It turns out this count is the same as that of the "labeled rooted forests." We prove in our Theorem 3.3 that the number of nilpotent k-file placements is equal to $\binom{n-1}{k}n^k$. This numerology is the first step towards understanding the S_n-module structure on the set of all nilpotent file placements. Towards this goal, we devote whole Section 4 to study the $n = 3$ case in detail.

At the beginning of Section 5 we show that the conjugation action on a nilpotent file placement does not alter the underlying (unlabeled) rooted forest. Moreover, we observe that this action is transitive on the labels. Let σ be a labeled rooted forest. We denote the resulting representation, that is to say, the orbit of σ under the conjugation action, by $o(\sigma)$ and call it the odun of σ. We observe in Theorem 5.4 that the representation $o = o(\sigma)$ is related by a simple operation to $o(\tau)$ if τ is the labeled rooted forest that is obtained from σ by removing the root. As a simple consequence of this fact, we obtain our first recursive relation among the characters of oduns of forests. Another important observation we make in the same section is that the adding of k new isolated vertices to a given rooted forest corresponds to tensoring the original representation by the k-dimensional standard representation of S_k. This is our Theorem 5.7. There is an even more general statement that we record in Remark 5.8: If the rooted forest σ is written as a disjoint union $\tau \cup \nu$ of two other rooted forests which do not have any identical rooted subtrees, then $o(\sigma) = o(\tau) \otimes o(\nu)$.

In Section 6 we prove our master plethysm result, Theorem 6.1, which states that if a rooted forest τ (on mk vertices) is comprised of k copies of the same rooted tree σ (on m vertices), then the odun of τ is given by the compositional product of the standard k-dimensional representation of S_k with $o(\sigma)$. Combined with Remark 5.8 this result gives us a satisfactorily complete description of the rooted forest representations.

The main purpose of Section 7 is to determine when the sign representation occurs in a given odun. The surprising combinatorial result of this section states that the sign representation occurs in $o(\sigma)$ if and only if σ is "blossoming." A rooted forest is called *blossoming* if it has no rooted subtree having (at least) two identical maximal terminal branches of odd length emanating from the same vertex. By counting blossoming forests, we show that the total number of occurrences of the sign representation in all rooted forest representations on n vertices is equal to 2^{n-2}.

The labeled rooted forest associated with a non-attacking rook placement has a distinguishing feature; it is a union of chains. From representation theory point of view the ordering of the chains does not matter, therefore, there is a correspondence between the number of rooted forest representations on non-attacking rooks and partitions. In Section 8 we make this precise. These results are simple applications of the previous sections.

In our "Final Remarks" section we give a formula for the dimension of a forest representation and compare our result with Knuth's hook-length formula. Finally, we close our paper in Section 10 by presenting tables of irreducible constituents of the nilpotent k file placements.

Acknowledgements. We are grateful to the referee whose careful reading and remarks improved the quality of our paper. We thank Michael Joyce, Brian Miceli, Jeff Remmel, and Lex Renner.

2 Preliminaries

2.1 Terminology of forests

A *rooted tree* σ is a finite collection of vertices such that there exists a designated vertex, called the root, and the remaining vertices are partitioned into a finite set of disjoint non-empty subsets $\sigma_1, \dots, \sigma_m$, each of which is a tree itself. We depict a tree by putting its root at the top so that the following terminology is logical: if a vertex a is connected by an edge to another vertex b that is directly above a, then a is called a *child* of b. A collection of rooted trees is called a rooted forest. For notational ease, unless it is necessary to stress its existence, we will omit the word "root" from our notation.

The elements of the set $[n] = \{1, \dots, n\}$ are often referred to as *labels*. A well known theorem of Cayley states that there are n^{n-2} labeled trees on n vertices. Since attaching a vertex (as the root) to a forest on n vertices gives a tree on $n+1$ vertices, Cayley's theorem is equivalent to the statement that there are $(n+1)^{n-1}$ labeled forests on n vertices.

2.2 Basic character theory

It is well-known that the irreducible representations of S_n are indexed by partitions of n. If λ is a partition of n (so we write $\lambda \vdash n$), then the corresponding irreducible representation is denoted by V_λ.

Let R^0 denote the ring of integers $\mathbb{Z}$ and let R^n, $n = 1, 2, \ldots$ denote the $\mathbb{Z}$-module spanned by irreducible characters of S_n. We set $R = \bigoplus_{n\geq 0} R^n$. In a similar fashion, let Λ denote the direct sum $\bigoplus_{n\geq 0} \Lambda^n$, where Λ^n is the $\mathbb{Z}$-module spanned by homogenous symmetric functions of degree n, and $\Lambda^0 = \mathbb{Z}$. Both of these $\mathbb{Z}$-modules are in fact $\mathbb{Z}$-algebras, and the Frobenius character map, which is defined by

$$\mathrm{ch}: \ R \ni \chi \mapsto \sum_{\rho\vdash n} z_\rho^{-1} \chi_\rho p_\rho \in \Lambda$$

is a $\mathbb{Z}$-algebra isomorphism. Here, z_ρ denotes the quantity $\prod_{i\geq 1} i^{m_i} m_i!$, where m_i is the number of occurrence of i as a part of ρ, χ_ρ is the value of the character χ on the conjugacy class indexed by the partition $\rho = (\rho_1, \ldots, \rho_r)$, and p_ρ stands for the power-sum symmetric function $p_\rho = p_{\rho_1} \cdots p_{\rho_r}$. Under Frobenius character map ch the irreducible character χ^λ of the representation V_λ is mapped to the Schur function $s_\lambda := \sum_{\rho\vdash n} z_\rho^{-1} \chi_\rho^\lambda p_\rho$. (This can be taken as the definition of a Schur function.) The monomial symmetric function associated with partition $\lambda = (\lambda_1, \ldots, \lambda_l)$ is defined as the sum of all monomials of the form $x_1^{\beta_1} \cdots x_l^{\beta_l}$, where $(\beta_1, \ldots, \beta_l)$ ranges over all distinct permutations of $(\lambda_1, \ldots, \lambda_l)$. Any of the sets $\{s_\lambda\}_{\lambda\vdash n}, \{m_\lambda\}_{\lambda\vdash n}$, and $\{p_\lambda\}_{\lambda\vdash n}$ forms a $\mathbb{Q}$-vector space basis for the vector space $\Lambda^n \otimes_{\mathbb{Z}} \mathbb{Q}$. Kostka numbers $K_{\lambda\mu}$ are defined as the coefficients in the expansion $s_\lambda = \sum_{\mu\vdash n} K_{\lambda\mu} m_\mu$.

2.3 Symmetric functions and plethysm

The plethysm of the Schur functions $s_\lambda \circ s_\mu$ is the symmetric function obtained from s_λ by substituting the monomials of s_μ for the variables of s_λ. To spell this out more precisely we follow [20]. The plethysm operator on symmetric functions is the unique map $\circ : \ \Lambda \times \Lambda \to \Lambda$ satisfying the following three axioms:

P1. For all $m, n \geq 1$, $p_m \circ p_n = p_{mn}$.

P2. For all $m \geq 1$, the map $g \mapsto p_m \circ g$, $g \in \Lambda$ defines a $\mathbb{Q}$-algebra homomorphism on Λ.

P3. For all $g \in \Lambda$, the map $h \mapsto h \circ g$, $h \in \Lambda$ defines a $\mathbb{Q}$-algebra homomorphism on Λ.

In general, computing the plethysm of two arbitrary symmetric functions is not easy. Fortunately, there are some useful formulas involving Schur functions:

$$s_\lambda \circ (g + h) = \sum_{\mu,\nu} c_{\mu,\nu}^\lambda (s_\mu \circ g)(s_\nu \circ h), \tag{2.1}$$

and

$$s_\lambda \circ (gh) = \sum_{\mu,\nu} \gamma_{\mu,\nu}^\lambda (s_\mu \circ g)(s_\nu \circ h). \tag{2.2}$$

Here, g and h are arbitrary symmetric functions, $c_{\mu,\nu}^\lambda$ is a scalar, and $\gamma_{\mu,\nu}^\lambda$ is $\frac{1}{n!}\langle \chi^\lambda, \chi^\mu \chi^\nu \rangle$. (The pairing $\langle \, , \rangle$ is the standard Hermitian inner product on the vector space of class functions on S_n.) In (2.1) the summation is over all pairs of partitions $\mu, \nu \subset \lambda$, and the summation in (2.2) is over all pairs of partitions μ, ν such that $|\mu| = |\nu| = |\lambda|$. In the special case when $\lambda = (n)$, or (1^n) we have

$$s_{(n)} \circ (gh) = \sum_{\lambda \vdash n} (s_\lambda \circ g)(s_\lambda \circ h), \tag{2.3}$$

$$s_{(1^n)} \circ (gh) = \sum_{\lambda \vdash n} (s_\lambda \circ g)(s_{\lambda'} \circ h), \tag{2.4}$$

where λ' denotes the conjugate of λ.

In a similar vein, if ρ denotes a partition, then the coefficient of s_μ in $p_\rho \circ h_n$ is given by $K^{(\rho)}_{\mu,n\rho}$, *the generalized Kostka number.* Since we need this quantity in one of our calculations, we define it. A generalized tableau of type ρ shape μ and weight $n\rho = (n\rho_1, n\rho_2, \dots, n\rho_m)$ is a sequence $T = (\nu^{(0)}, \nu^{(1)}, \dots, \nu^{(m)})$ of partitions satisfying the following conditions

1. $0 = \nu^{(0)} \subset \nu^{(1)} \subset \cdots \subset \nu^{(m)} = \mu$;
2. $|\nu^{(j)} - \nu^{(j-1)}| = n\rho_j$ for $1 \le j \le m$;
3. $\nu^{(j)} \approx_{\rho_j} \nu^{(j-1)}$ for $1 \le j \le m$. (See [21] Chapter I, §5, Example 24 for definition of $\approx_{\rho_j}$.)

For such tableau, define $\sigma(T) := \prod_{j=1}^m \sigma_{\rho_j}(\nu^{(j)}/\nu^{(j-1)})$ which is ± 1. (Once again, see [21] Chapter I, §5, Example 24 for definition of $\sigma_{\rho_j}(\cdot)$). Finally, we define $K^{(\rho)}_{\mu,n\rho}$ as the sum of $\sigma(T)$'s

$$K^{(\rho)}_{\mu,n\rho} := \sum_T \sigma(T),$$

where the sum ranges over all generalized tableau T of type ρ, shape μ, and of weight $n\rho$.

3 Re-counting nilpotent file placements

The relationship between file placements and directed acyclic graphs is briefly mentioned in the introduction section; the map φ that sends a labeled forest on n vertices (labeled with $\{1, \dots, n\}$) to its incidence matrix is a bijection onto its image. Furthermore, the image of φ is precisely the set of $n \times n$ 0/1 matrices with at most one 1 in each column. Our goal in this section is to count, in the image of φ, the number of labeled forests.

Let us call a file placement f nilpotent if its matrix (1.1) corresponds to a digraph without cycles or loops. We start with a simple lemma to justify the nomenclature.

Lemma 3.1. *A file placement is nilpotent if and only if its associated matrix (as defined in Section 1) is nilpotent.*

Proof. Let $f : A \to [n]$ denote the partial transformation representing a file placement on $[n] \times [n]$. If the graph of f has a non-trivial cycle, then there exists a sequence $i_1, \dots, i_r$ numbers from the domain A of f such that $f(i_1) = i_2, f(i_2) = i_3, \dots, f(i_r) = i_1$. It follows that for any $m > 0$, $f^m(i_1) = f^{m \mod r}(i_1) \in \{i_1, \dots, i_r\}$, hence no power of f can be zero. Conversely, if f is nilpotent, then for any $a \in A$, some power of f vanishes on a. Therefore, $a, f(a), f^2(a), \dots$ does not return to a to become a cycle. □

Remark 3.2. The proof of the above lemma implies that a file placement $f : A \to [n]$ is nilpotent if and only if there does *not* exist a subset $D \subset A$ such that $f(D) = D$. This observation is recorded in [18].

Theorem 3.3. *The number of nilpotent k-file placements is equal to $\binom{n-1}{k} n^k$.*

Proof. It is a well known variation of the Cayley's theorem that the number of labeled forests on n vertices with k roots is equal to $\binom{n-1}{k-1} n^{n-k}$. See [4], Theorem D, pg 70. Since the labeled directed graph of a nilpotent partial transformation has no cycles, it is a disjoint union of trees and the total number of vertices is n. Therefore, it remains to show that the labeled forest of a k-file placement f has exactly $n - k$ connected components. We prove this by induction on k.

If $k = 1$, then the forest of f has one component on two vertices and $n - 2$ singletons. Therefore, the base case is clear. Now we assume that our claim is true for $k - 1$ and prove it for k-file placements. Let $j \in A$ be a number that is not contained in the image of f whose existence is guaranteed by Remark 3.2. Define $\tilde{f} : A - \{j\} \to [n]$ by setting $\tilde{f}(i) = f(i)$ for $i \in A - \{j\}$. Therefore, $\tilde{f}$ is a $(k - 1)$-file placement agreeing with f at all places except at $\{j\}$, where it is undefined. By our induction hypothesis, the forest of $\tilde{f}$ has exactly $n - (k - 1)$ connected components. Observe that the forest of f differs from that of $\tilde{f}$ by exactly one directed edge from j to $f(j)$. Since $\{j\}$ is a connected component of $\tilde{f}$, the number of connected components of f is one less than that of $\tilde{f}$, hence the proof is complete. □

Remark 3.4. Theorem 3.3 gives the number of partial transformations on $[n]$ having exactly k elements in their domains. There is a related result of Laradji and Umar that appeared in [18]. For completeness of the section let us briefly present it: Let $N(J_r)$ denote the set of partial transformations $f : A \to [n]$ with $|f(A)| = r$. In their Theorem 3, Laradji and Umar compute that

$$|N(J_r)| = \binom{n}{r} S(n, r + 1) r!,$$

where $S(n, r + 1)$ is the Stirling number of the second kind, namely the number of set partitions of $[n]$ into $r + 1$ non-empty blocks.

Remark 3.5. Let $N_{k,n}$ denote the number of nilpotent k-file placements on $[n]$ and define $E_x(y)$ by

$$E_x(y) := \sum_{n \geq 1} \sum_{k \geq 0}^{n-1} N_{k,n} x^k \frac{y^n}{n!} = \sum_{n \geq 1} (1 + nx)^{n-1} \frac{y^n}{n!}. \tag{3.6}$$

For the following identities, see [10, Chapter 5].

1. $E^{-x}\ln E = y$,
2. $E^{\alpha} = \sum_{n\geq 0} \alpha(\alpha+nx)^{n-1}\frac{y^n}{n!}$ for all $\alpha \in \mathbb{R}$,
3. $\frac{E^{\alpha}}{1-xyE^{x}} = \sum_{n\geq 0}(\alpha+nx)^{n}\frac{y^n}{n!}$ for all $\alpha \in \mathbb{R}$.

Now we consider unlabeled trees. Let t_{n+1} denote the number of rooted trees, up to isomorphism, on $n+1$ vertices. Manipulations with generating functions lead to the following non-trivial recurrence for t_n's:

$$t_{n+1} = \frac{1}{n}\sum_{k=1}^{n}\left(\sum_{d|k} dt_d\right)t_{n-k+1} \qquad (t_1 = 1). \tag{3.7}$$

In contrast with the labeled case, eqn. (3.7) indicates that the number of unlabeled forests is not as easily expressible as one wishes.

4 A case study

We denote by $\mathcal{C}_{k,n}$ the set of nilpotent k-file placements on $[n]$ and set

$$\mathcal{C}_n := \bigcup_{k=0}^{n-1}\mathcal{C}_{k,n}.$$

Remark 4.1. It follows from the proof of Theorem 3.3 that the elements of $\mathcal{C}_{k,n}$ are in bijection with the labeled rooted forests on n vertices with $n-k$ connected components.

Obviously, there is a single nilpotent partial transformation on $\{1\}$. For $n=2$, the nilpotent partial transformations are

$$\mathcal{C}_{0,2} = \{[0,0]\}, \qquad \mathcal{C}_{1,2} = \{[0,1],[2,0]\},$$

and for $n=3$ we have

$$\begin{aligned}
&\mathcal{C}_{0,3} = \{[0,0,0]\},\\
&\mathcal{C}_{1,3} = \{[0,1,0],[0,3,0],[0,0,1],[0,0,2],[2,0,0],[3,0,0]\},\\
&\mathcal{C}_{2,3} = \{[0,1,2],[0,1,1],[0,3,1],[2,0,1],[2,0,2],[2,3,0],[3,0,2],[3,1,0],[3,3,0]\}.
\end{aligned}$$

The symmetric group S_3 acts on each of the sets $C_{i,3}$, $i=0,1,2$ by conjugation. The table of corresponding character values are easily determined by counting fixed points of the action. The first column in Table 1 is the list of representatives for each conjugacy class in S_3. In the next table we have the character values of all irreducible representations of S_3. In the last column, we have listed the sizes of the corresponding conjugacy classes: Now, from Tables 1 and 2 we see that

g	$\mathcal{C}_{0,3}$	$\mathcal{C}_{1,3}$	$\mathcal{C}_{2,3}$
(1)(2)(3)	1	6	9
(12)(3)	1	0	1
(123)	1	0	0

Table 1: The character values of the S_3-representations on nilpotent 3-file placements

$g \in S_3$	$V_{(3)}$	$V_{(2,1)}$	$V_{(1,1,1)}$	$c(g)$
(1)(2)(3)	1	2	1	1
(12)(3)	1	0	-1	3
(123)	1	-1	1	2

Table 2: Irreducible character values of S_3

- $\mathcal{C}_{0,3}$ is the trivial representation $V_{(3)}$.
- $\mathcal{C}_{1,3}$ is equal to $V_{(3)} \oplus V_{(2,1)}^2 \oplus V_{(1,1,1)}$.
- $\mathcal{C}_{2,3} = V_{(3)}^2 \oplus V_{(2,1)}^3 \oplus V_{(1,1,1)}$.

Going back to cases where $n = 1$ and $n = 2$ we compute also that, as a representation of S_1, $\mathcal{C}_{0,1}$ is the unique irreducible (trivial) representation $V_{(1)}$ of S_1. Similarly, $\mathcal{C}_{0,2} = V_{(2)}$, and $\mathcal{C}_{1,2} = V_{(2)} \oplus V_{(1,1)}$. We listed the decomposition tables for $n = 4, 5$ and $n = 6$ at the end of the paper.

5 Induced representations

Let n be a positive integer and denote by $\chi^n := \chi^{\mathcal{C}_n}$ the character of the conjugation action on nilpotent file placements. This notation (χ^n) should not be confused with the character $\chi^{(n)}$ of the irreducible representation associated with the partition $\lambda = (n)$. Note that the conjugation action does not change the number of elements in the domain of a partial function, therefore, $\mathcal{C}_{k,n}$ is closed under the S_n-action (see Remark 3.4). In other words, $\mathcal{C}_{k,n}$ defines an S_n-submodule of the representation on nilpotent file placements. We denote by $\chi^{k,n} := \chi^{\mathcal{C}_{k,n}}$ the character of the conjugation action on nilpotent k-file placements. Clearly, $\chi^n = \sum_{k=0}^{n-1} \chi^{k,n}$. It is also clear that if $k = 0$, then the S_n-module defined by $\mathcal{C}_{0,n}$ is the trivial 1 dimensional representation, so we will focus on the cases where $k \geq 1$.

Lemma 5.1. *For $i = 1, \dots, n-1$, let $(i, i+1) \in S_n$ denote the simple transposition that interchanges i and $i+1$. If $\sigma = \sigma_f$ is a labeled tree corresponding to a nilpotent k-file placement f, then the underlying (unlabeled) tree of $(i, i+1) \cdot \sigma$ is equal to that of σ. Moreover, if o denotes an unlabeled tree on n vertices, then S_n acts transitively on the set of elements $\sigma \in \mathcal{C}_{k,n}$ whose underlying tree is equal to o.*

Proof. Let $[f_1, f_2, \dots, f_n]$ denote the one-line notation for f. The conjugation action of $(i, i+1)$ on f has the following effect: 1) the entries f_i and f_{i+1} are interchanged, 2) if $f_l = i$ for some $1 \leq l \leq n$, then f_l is replaced by $i+1$. Similarly, if $f_{l'} = i+1$ for some $1 \leq l' \leq n$, then $f_{l'}$ is replaced by i. It is clear that these operations do not change the underlying graph structure, they act as permutations on labels only.

To prove the last statement we fix a labeled tree σ. It is enough to show the existence of a permutation which interchanges two chosen labels i and k on σ without changing any other labels. Looking at the one-line notation for f, we see that the action of transposition (i, k) gives the desired result. □

Caution: Recall our terminology from the introductory section; the odun $o(\sigma)$ of σ is the S_n-representation on the orbit $S_n \cdot \sigma$. By Lemma 5.1, we see that the odun is completely determined by the underlying (unlabeled) forest. Therefore, it will cause no harm to use the word "odun" for the underlying unlabeled structure as well.

Corollary 5.2. *The multiplicity of the trivial character in χ^n is the number of rooted forests on n vertices. Equivalently, $\langle \chi^{(n)}, \chi^n \rangle$ = number of trees on $n+1$ vertices, which is equal to t_{n+1} of (3.7).*

Proof. This is a standard fact: The multiplicity of the trivial representation in any permutation representation is equal to the number of orbits of the action. By Lemma 5.1, this number is equal to the number of oduns (unlabeled forests) on n vertices. □

Example 5.3. Let σ be a labeled forest and let $o = o(\sigma)$ denote its odun. We denote the corresponding character by $\chi^{o(\sigma)}$. Here, we produce three examples of forest representations, decomposed into irreducible constituents. We will use this observation in the sequel.

1. $o(\sigma) =$  $\Longrightarrow$ $\chi^{o(\sigma)} = \chi^{(3)} + 2\chi^{(2,1)} + \chi^{(1^3)}$

2. $o(\sigma) =$ 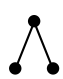 $\Longrightarrow$ $\chi^{o(\sigma)} = \chi^{(2,1)} + \chi^{(1^3)}$

3. $o(\sigma) =$ 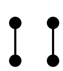 $\Longrightarrow$ $\chi^{o(\sigma)} = \chi^{(4)} + \chi^{(3,1)} + 2\chi^{(2,2)} + \chi^{(2,1,1)} + \chi^{(1^4)}$

Theorem 5.4. *Let σ be a labeled rooted forest corresponding to an element f of $\mathcal{C}_{k,n}$ and let $\chi^{\mathbb{1}}$ denote the character of the unique (1 dimensional) representation of S_1. If $\widetilde{\sigma}$ denotes the labeled tree obtained from σ by connecting all of its roots to a new vertex, which is labeled by $n+1$, then*

$$\chi^{o(\widetilde{\sigma})} = \chi^{\mathbb{1}} \cdot \chi^{o(\sigma)}.$$

Proof. We will compute the stabilizer subgroup of $\widetilde{\sigma}$, the labeled tree that is obtained from σ by adding a vertex with label $n+1$ as a root. We observed earlier that the conjugation action of the symmetric group is transitive on the labels. Looking at the label of the new vertex, we see that the S_{n+1} orbit of $\widetilde{\sigma}$ essentially breaks down into $n+1$ subsets. The stabilizer subgroup of any element in any of these subsets is isomorphic to $S_1 \times \mathrm{Stab}_{S_n}(\sigma)$. Therefore,

$$o(\widetilde{\sigma}) = \oplus_{\pi \in S_{n+1}/S_n} \pi \cdot o(\sigma). \tag{5.5}$$

But notice that by (5.5) we are repeating the definition of an induced representation of $\mathbb{C}1 \otimes o(\sigma)$ from $S_1 \times S_n$ to S_{n+1}. Hence, the proof is complete. □

Corollary 5.6. *The character $\chi^{n-1,n}$ of $\mathcal{C}_{n-1,n}$ is equal to $\chi^{\mathbb{1}}(\chi^{0,n-1}+\chi^{1,n-1}+\cdots+\chi^{n-2,n-1})$.*

Proof. By adding a root we go from rooted forests to trees. We state this in terms of nilpotent partial transformations; an element of $\mathcal{C}_{n-1,n}$ corresponds to a labeled tree on n vertices and by removing the root (which can be assumed to have the label n by the transitivity of the S_n-action) we obtain a labeled rooted forest that is represented by a nilpotent partial function in $\mathcal{C}_{n-1,n}$ (see Remark 4.1). The rest of the proof follows from Theorem 5.4. □

The idea of the proof of our next result is identical to that of Theorem 5.4, so we omit it.

Theorem 5.7. *Let σ be a labeled forest with m connected components each of which has at least two vertices. If $\widetilde{\sigma}$ is the labeled forest obtained from σ by adding k isolated roots (hence it has $m+k$ connected components), then*

$$\chi^{o(\widetilde{\sigma})} = \chi^{(k)} \cdot \chi^{o(\sigma)}.$$

Remark 5.8. More general than Theorem 5.4 with an almost identical proof is the following statement: Suppose ν is a labeled forest of the form $\tau \cup \sigma$ (disjoint union), where the underlying trees of τ and σ are nonidentical. In this case,

$$\chi^{\nu} = \chi^{\tau} \cdot \chi^{\sigma}. \tag{5.9}$$

Using Theorems 5.7 and 5.4 we will perform a sample calculation of the characters for small k. If W is a representation and r is a positive integer, then by W^r we mean the representation on the direct sum of r copies of W.

Proposition 5.10. *Let n be a positive integer. If $n \geq 3$, then we have*

(i) $\mathcal{C}_{1,n} = V_{(n)} \oplus V^2_{(n-1,1)} \oplus V_{(n-2,2)} \oplus V_{(n-2,1,1)}$;

(ii) $\mathcal{C}_{2,n} = (V_{(3)} \oplus V^2_{(2,1)} \oplus V_{(1^3)}) \otimes V_{(n-3)} \oplus (V^2_{(2,1)} \oplus V_{(1^3)}) \otimes V_{(n-3)} \oplus (V_{(4)} \oplus V_{(3,1)} \oplus V^2_{(2,2)} \oplus V_{(2,1,1)} \oplus V_{(1^4)}) \otimes V_{(n-4)}$.

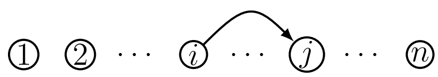

Figure 5.1: The labeled rooted forest form of a typical element of $\mathcal{C}_{1,n}$.

Proof. (i) An element f from $\mathcal{C}_{1,n}$ is represented in one-line notation by a sequence $[f_1, \ldots, f_n]$ of length n having a single nonzero entry, say at the i-th position. Furthermore, if $f_i = j$, then we have $j \in [n] - \{i\}$. Thus, the corresponding labeled tree $\sigma = \sigma_f$ is an array of n vertices labeled from 1 to n, and the i-th vertex is connected to the j-th by a directed edge. See Figure 5.1. The rest of the proof of (i) follows from Theorem 5.7.

(ii) It is easy to verify that the underlying forest of an element $\sigma \in \mathcal{C}_{2,n}$ is one of the three forests which are depicted in Figure 5.2. Thus, the proof follows from Theorem 5.7 in

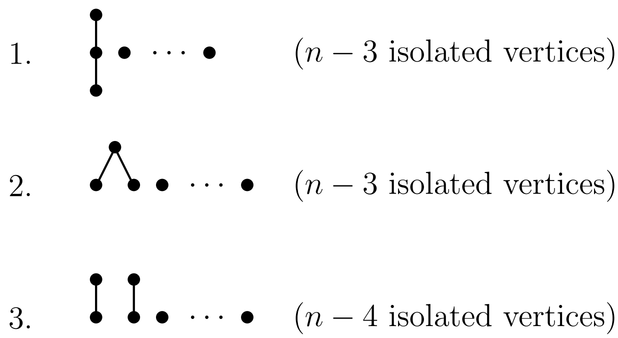

Figure 5.2: The list of possible oduns in $\mathcal{C}_{2,n}$.

view of Example 5.3. □

6 Plethysm and labeled trees

It follows from Remark 5.8 that if a labeled tree σ is obtained from a collection of mutually nonidentical trees $\sigma_1, \ldots, \sigma_r$ by adding a root, then the character of $o(\sigma)$ is given by $\chi^{o(\sigma)} = \chi^{\mathbb{1}} \prod_{i=1}^{r} \chi^{o(\sigma_i)}$. Therefore, as far as the general case concerned, we need to focus on the following (remaining) situation: the underlying forest of σ is comprised of r copies of the same tree τ.

Theorem 6.1. *Let σ be a labeled forest with r connected components all of which have the same underlying tree, τ. In this case, the representation $o(\sigma)$ is equal to the plethysm of $\chi^{(r)}$ with $\chi^{o(\tau)}$. More precisely,*

$$\chi^{o(\sigma)} = \chi^{(r)} \circ \chi^{o(\tau)}.$$

Proof. Let m denote the number of vertices of τ. Then σ has $n := mr$ vertices. The action of S_n on σ has the following defining property: if i and j are the labels of two vertices connected by an edge, then for any permutation $g \in S_n$ the vertices with labels $g(i)$ and $g(j)$ have an edge between them as well. In other words, the action of S_n maps the set of all labels on a connected component to the labels on another connected component. (Of course, these two connected components might be the same.) It follows that the stabilizer subgroup in S_n of σ is of the form $S_r \wr Stab_{S_m}(\tau)$ and our representation $o(\sigma)$ is isomorphic to $\mathrm{Ind}_{S_r \wr Stab_{S_m}(\tau)}^{S_n} 1$. It is well-known that the character of such a representation is given by the plethysm of the corresponding characters. (See Macdonald [21] Appendix A.) Since the character of permutation representation of S_r on r letters is $\chi^{(r)}$, the proof is complete. □

Corollary 6.2. *Let $\lambda = (\lambda_1, \dots, \lambda_r)$ be a partition and let σ be a labeled forest comprised of λ_1 copies of the (unlabeled) tree σ_1, λ_2 copies of (unlabeled) tree σ_2, and so on. If the underlying trees of σ_i's ($i = 1, 2, \dots$) are mutually nonidentical, then the character of $o(\sigma)$ is*

$$\chi^{o(\sigma)} = (\chi^{(\lambda_1)} \circ \chi^{o(\sigma_1)}) \cdot (\chi^{(\lambda_2)} \circ \chi^{o(\sigma_2)}) \cdot \dots \cdot (\chi^{(\lambda_r)} \circ \chi^{o(\sigma_r)}). \tag{6.3}$$

Proof. The proof follows from Remark 5.8 and Theorem 6.1. □

Corollary 6.4. *The character of the k-file placements $\mathcal{C}_{k,n}$ is given by*

$$\chi^{k,n} = \sum_{(m_1,\dots,m_r)\vdash n-k} \sum_{(o_1,\dots,o_r)} (\chi^{(m_1)} \circ \chi^{o_1}) \cdot (\chi^{(m_2)} \circ \chi^{o_2}) \cdot \dots \cdot (\chi^{(m_r)} \circ \chi^{o_r}),$$

where the second summation is over all tuples of nonidentical trees such that $\sum m_i |o_i| = n$.

Proof. This is an immediate consequence of Corollary 6.2. □

7 The sign representation

In this section we compute the multiplicity of the sign representation in an odun. Surprisingly, asymmetry in the underlying rooted tree is a source of regularity in the associated representation.

Let $o = o(\sigma)$ be the odun of a labeled tree and let F_o denote the corresponding Frobenius character. Removing the root from σ gives a forest τ. Let $\sigma_1, \dots, \sigma_r$ denote the list of distinct connected components (subtrees) of τ and let $o_1, \dots, o_r$ denote the corresponding list of oduns. By Corollary 6.2 we see that

$$F_o = s_{(1)} \cdot s_{(i_1)} \circ F_{o_1} \cdot s_{(i_2)} \circ F_{o_2} \cdot \dots \cdot s_{(i_r)} \circ F_{o_r}, \tag{7.1}$$

where $i_1, \dots, i_r$ are the multiplicities in σ of the subtrees $\sigma_1, \dots, \sigma_r$ in the listed order. It is easy to see by using Frobenius character map that the multiplicity of the sign representation in F_o is equal to $\langle F_o, s_{(1^m)} \rangle$, where m is the number of vertices in σ. For $j = 1, \dots, r$, let k_j denote i_j times the number of vertices in σ_j. Our basic observation is that $s_{(1^m)}$ can occur in F_o only if $s_{(1^{k_j})}$ occurs in $s_{(i_j)} \circ F_{o_j}$ for $j = 1, \dots, r$. Therefore, we will focus initially on the computation of the multiplicity of $s_{(1^n)}$ in a plethysm of the form $s_{(k)} \circ F_o$, where n and k are positive integers, o is the odun of a labeled tree on m vertices with $n = km$. For this goal, let us start with an even more general expression: $\langle s_{(1^m)}, s_\lambda \circ F_o \rangle$, where λ is a partition of k.

Let h denote the Frobenius character of τ, which is equal to $F_o/s_{(1)} = \prod_{j=1}^r s_{(i_j)} \circ F_{o_j}$. Then

$$\begin{aligned} s_\lambda \circ F_o &= s_\lambda \circ (s_{(1)}h) && (7.2) \\ &= \sum_{\mu,\nu} \gamma^\lambda_{\mu,\nu} (s_\mu \circ s_{(1)})(s_\nu \circ h) \qquad \text{(by (2.2))} \\ &= \sum_{\mu,\nu} \gamma^\lambda_{\mu,\nu} s_\mu (s_\nu \circ h) \qquad \text{(by the axioms of plethysms).} && (7.3) \end{aligned}$$

Recall that $\gamma^\lambda_{\mu,\nu} = \frac{1}{k!}\langle \chi^\lambda, \chi^\mu \chi^\nu \rangle$. On one hand, whenever $\lambda = (k),\ \mu = (1^k)$ we have

$$\begin{aligned} \gamma^{(k)}_{(1^k),\nu} &= \frac{1}{k!} \sum_{w \in S_k} \chi^{(k)}(w) \chi^{(1^k)}(w) \chi^\nu(w) \\ &= \frac{1}{k!} \sum_{w \in S_k} \chi^{(1^k)}(w) \chi^\nu(w) \qquad \text{(since } \chi^{(k)}(w) = 1 \text{ for all } w \in S_k) \\ &= \langle \chi^{(1^k)}, \chi^\nu \rangle \\ &= \begin{cases} 0 & \text{if } \nu \neq (1^k); \\ 1 & \text{if } \nu = (1^k). \end{cases} \end{aligned}$$

By the same token, but more generally we have

$$s_{(k)} \circ F_o = \sum_\mu s_\mu (s_\mu \circ h). \tag{7.4}$$

Thus, in (7.4) the sign representation $s_{(1^m)}$ occurs if and only if $\mu = (1^k)$ and $s_{(1^{m-k})}$ is a summand of $s_\nu \circ h$. Therefore,

$$\langle s_\lambda \circ F_o, s_{(1^m)} \rangle = \langle s_{(1^{m-k})}, \sum_\nu \gamma^\lambda_{(1^k),\nu} s_\nu \circ h \rangle. \tag{7.5}$$

On the other hand, $\lambda = (1^k)$ implies that

$$\begin{aligned} \gamma^{(1^k)}_{(1^k),\nu} &= \frac{1}{k!} \sum_{w \in S_k} \chi^{(1^k)}(w) \chi^{(1^k)}(w) \chi^\nu(w) \\ &= \frac{1}{k!} \sum_{w \in S_k} \chi^\nu(w) \qquad \text{(since } \chi^{(1^k)}(w) = \pm 1 \text{ for all } w \in S_k) \\ &= \begin{cases} 0 & \text{if } \nu \neq (k); \\ 1 & \text{if } \nu = (k). \end{cases} \end{aligned}$$

More generally, by using the same idea we obtain $s_{(1^k)} \circ F_o = \sum_{\mu \vdash k} s_{\mu'}(s_\mu \circ h)$, where k is a positive integer. The reasoning which we used right after equation (7.3) gives more:

$$\begin{aligned}
\langle s_{(1^m)}, s_\lambda \circ F_o \rangle &= \langle s_{(1^m)}, \sum_{\mu,\nu} \gamma^\lambda_{\mu,\nu} s_\mu (s_\nu \circ h) \rangle \\
&= \langle s_{(1^m)}, \sum_{\nu} \gamma^\lambda_{(1^k),\nu} s_{(1^k)} (s_\nu \circ h) \rangle \qquad (k \text{ is equal to } |\lambda|) \\
&= \langle s_{(1^m)}, \sum_{\nu} \gamma^{(1^k)}_{\lambda,\nu} s_{(1^k)} (s_\nu \circ h) \rangle \\
&= \langle s_{(1^m)}, s_{(1^k)} (s_{\lambda'} \circ h) \rangle \\
&= \langle s_{(1^{m-k})}, s_{\lambda'} \circ h \rangle
\end{aligned}$$

In conclusion, we have the following 'simplification/duality' result:

Lemma 7.6. *Let F_o denote the Frobenius character of the odun o of a tree on n vertices, and let $h = F_\tau$ denote the Frobenius character of the odun of the rooted forest on $n-1$ vertices obtained from o by removing its root. If k and m are two positive integers, then we the following identities hold true:*

1. $s_{(k)} \circ F_o = \sum_{\mu \vdash k} s_\mu (s_\mu \circ h)$;
2. $s_{(1^k)} \circ F_o = \sum_{\mu \vdash k} s_{\mu'} (s_\mu \circ h)$;
3. $\langle s_{(1^m)}, s_\lambda \circ F_o \rangle = \langle s_{(1^{m-k})}, s_{\lambda'} \circ h \rangle$.

In particular, the following equations hold true:

3.1 $\langle s_{(1^m)}, s_{(k)} \circ F_o \rangle = \langle s_{(1^m)}, s_{(1^k)} (s_{(1^k)} \circ h) \rangle = \langle s_{(1^{m-k})}, s_{(1^k)} \circ F_\tau \rangle$;

3.2 $\langle s_{(1^m)}, s_{(1^k)} \circ F_o \rangle = \langle s_{(1^m)}, s_{(1^k)} (s_{(k)} \circ h) \rangle = \langle s_{(1^{m-k})}, s_{(k)} \circ F_\tau \rangle$.

As far as the sign representation is concerned, we have the following crucial definition.

Definition 7.7. A (rooted) subtree a of a rooted tree is called a *terminal branch* (TB for short) if any vertex of a has at most one child. A *maximal terminal branch* (or, MTB for short) is a terminal branch that is not a subtree of any terminal branch other than itself. The length (or height) of a TB is the number of vertices it has. We call a rooted tree *blossoming* if all of its MTB's are of even length, or no two odd length MTB's of the same length are connected to the same parent. A rooted tree which is not blossoming is called *dry*.

Lemma 7.8. *If an odun o is an MTB of length l, then for any $m \geq 2$*

$$\langle s_{(m)} \circ F_o, s_{(1^{ml})} \rangle = \begin{cases} 1 & \textit{if } l \textit{ is an even number;} \\ 0 & \textit{otherwise.} \end{cases}$$

Proof. Observe that removing the root from an MTB results in another MTB whose length is one less than the original's. The rest of the proof follows from Lemma 7.6 and axiomatic properties of plethysm. □

We proceed with extending our Definition 7.7 to forests.

Definition 7.9. We call a rooted forest blossoming if the rooted tree obtained by adding a new root to the forest is a blossoming tree. Otherwise, the forest is called dry.

Note that if a forest is blossoming, then all of its connected components are blossoming. Also, if a single connected component is dry, then the whole forest is dry. In Figure 7.1 we have listed all blossoming forests up to 4 vertices.

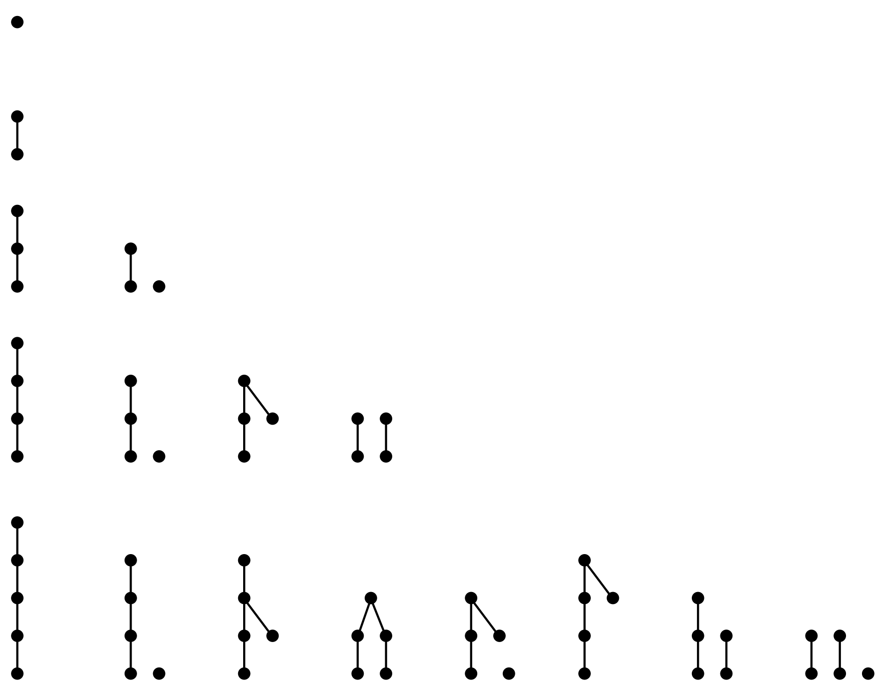

Figure 7.1: Blossoming forests on $1 \leq n \leq 5$ vertices.

Proposition 7.10. *Let λ be a partition with $|\lambda| = l > 1$ and let τ be a rooted forest on m vertices.*

1. *If τ is a blossoming forest, then* $\langle s_\lambda \circ F_\tau, s_{(1^{ml})}\rangle = \begin{cases} 1 & \text{if } \lambda = (1^l); \\ 0 & \text{otherwise.} \end{cases}$

2. *If τ is a dry forest, then $\langle s_\lambda \circ F_\tau, s_{(1^{ml})}\rangle = 0$.*

Proof. We prove both of our claims by induction on m. It is straightforward to verify them for $m = 1$ and 2, so, we assume that our claim is true for all forests with $|\tau| < m$. Suppose

$o_1, \ldots, o_r$ are the oduns of the connected components of the forest τ. Note that if τ is a rooted tree, then by part 3 of Lemma 7.6 our problem reduces to the forest case.

Let g_1 denote $s_{(i_1)} \circ F_{o_1}$ and let g_2 denote $\prod_{j=2}^{r} s_{(i_j)} \circ F_{o_j}$ so that $F_\tau = g_1 g_2$.

$$\begin{aligned} s_\lambda \circ F_\tau &= s_\lambda \circ (g_1 g_2) \\ &= \sum_{\mu,\mu} \gamma^\lambda_{\mu,\nu} (s_\mu \circ (s_{(i_1)} \circ F_{o_1}))(s_\nu \circ g_2) \\ &= \sum_{\mu,\nu} \gamma^\lambda_{\mu,\nu} \left(\sum_{\tilde{\mu}} \left(\sum_{\tilde{\nu}} z_{\tilde{\nu}}^{-1} \chi^\mu_{\tilde{\nu}} K^{(\tilde{\nu})}_{\tilde{\mu}, i_1 \tilde{\nu}} \right) (s_{\tilde{\mu}} \circ F_{o_1}) \right) (s_\nu \circ g_2). \end{aligned}$$

Now we are ready to start the induction argument. On one hand, if τ is dry, at least one of its rooted subtrees is dry. Without loss of generality, let o_1 denote the dry one. Thus, $\langle s_{\tilde{\mu}} \circ F_{o_1}, s_{(1^{|\tilde{\mu}||o_1|})} \rangle = 0$ for all partitions $\tilde{\mu}$, implying that $\langle s_\lambda \circ F_\tau, s_{(1^{ml})} \rangle = 0$. On the other hand, if τ is blossoming, all of its rooted subtrees are blossoming. Therefore,

$$\langle s_{\tilde{\mu}} \circ F_{o_1}, s_{(1^{|\tilde{\mu}||o_1|})} \rangle = \begin{cases} 1 & \text{if } \tilde{\mu} \text{ is of the form } (1^s),\ s = |\tilde{\mu}|; \\ 0 & \text{otherwise.} \end{cases}$$

In this case, that is when $\tilde{\mu} = (1^s)$, we calculate that $K^{(\tilde{\nu})}_{\tilde{\mu}, i_1 \tilde{\nu}} = 1$. (Follows from the explicit description of the generalized Kostka numbers as given in [5]). Therefore,

$$\sum_{\tilde{\nu}} z_{\tilde{\nu}}^{-1} \chi^\mu_{\tilde{\nu}} K^{(\tilde{\nu})}_{\tilde{\mu}, i_1 \tilde{\nu}} = \sum_{\tilde{\nu}} z_{\tilde{\nu}}^{-1} \chi^\mu_{\tilde{\nu}} = \begin{cases} 1 & \text{if } \mu = (s)\ ; \\ 0 & \text{otherwise.} \end{cases}$$

But it follows from $\mu = (s)$ that $\gamma^\lambda_{\mu,\nu} = \delta_{\lambda,\nu}$, the Kronecker delta function. Thus, $\langle \lambda s_{(1^m)}, s_\lambda \circ g_1 g_2 \rangle$ reduces to $\langle \lambda s_{(1^m)}, s_\lambda \circ g_2 \rangle$, which, by induction, is equal to 1 if τ, hence g_2 is blossoming. □

As an application of Proposition 7.10 we determine the multiplicity of the sign representation in $\mathcal{C}_{k,n}$. It boils down to the counting of blossoming trees.

Theorem 7.11. *In $\mathcal{C}_{n-1,n}$ the sign representation occurs exactly 2^{n-3} times, and in $\mathcal{C}_n$ the sign representation occurs 2^{n-2} times.*

By Corollary 5.6, it suffices to prove that $\langle \chi^n, \chi^{(1^n)} \rangle = 2^{n-2}$. We cast our problem in symmetric function language. Let T_n denote the Frobenius character $\mathrm{ch}(\chi^n)$, and let $T_{n-1,n}$ denote $\mathrm{ch}(\chi^{n-1,n})$. We already know that $T_{n-1,n} = s_{(1)} T_n$, hence that $\langle T_{n-1,n}, s_{(1^n)} \rangle = \langle T_{n-1}, s_{(1^{n-1})} \rangle$. Therefore, by Proposition 7.10, it suffices to find the number of blossoming trees on n vertices.

Proposition 7.12. *The number of blossoming forests on n vertices is 2^{n-2}.*

Proof. Let a_n denote the number of blossoming forests on n vertices without any isolated vertices, and let b_n denote the number of blossoming forests on n vertices with an isolated vertex. Set $d_n = a_n + b_n$. Clearly, d_n is the total number of blossoming forests on n vertices. Few values of a_n and b_n's are $a_1 = a_2 = a_3 = 1, a_4 = 3, a_5 = 5$, and $b_1 = b_2 = 0, b_3 = b_4 = 1, b_5 = 3$. See Figure 7.1.

There are obvious relations among a_n's, b_n's, and d_n's. For example, adding an isolated vertex to a blossoming forest without an isolated vertex gives

$$b_{n+1} = a_n \quad \text{for all } n \geq 1. \tag{7.13}$$

Similarly, we obtain all blossoming forests on n vertices with no isolated vertex by attaching a new single vertex to the isolated vertex of a forest, or by attaching a new root to all of the connected components. We depict this in Figure 7.2. The relation we obtain here is

$$a_n = b_{n-1} + d_{n-1} = 2b_{n-1} + a_{n-1} \quad \text{for } n \geq 3. \tag{7.14}$$

By combining (7.13) and (7.14) we arrive at a single recurrence,

$$a_n = 2a_{n-2} + a_{n-1} \tag{7.15}$$

with initial conditions $a_0 = a_1 = 0$, $a_2 = 1$. Let $f(x)$ denote the generating series $f(x) =$

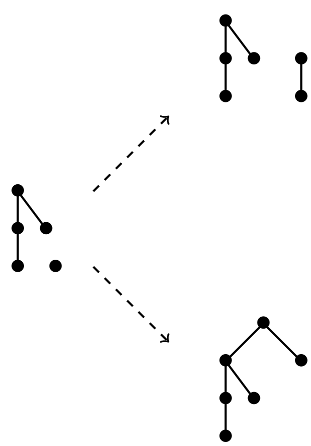

Figure 7.2: Adding a new root to blossoming forests.

$\sum_{n\geq 0} a_n x^n$. A straightforward generating function computation gives

$$f(x) = \frac{x^2}{1 - x - x^2}.$$

Denoting the generating series of d_n by $G(x)$, the relationship $d_n = a_n + b_n = a_n + a_{n-1}$ tells us

$$G(x) = f(x) + xf(x) = \frac{x^2 + x^3}{1 - x - x^2} = \frac{x^2}{1 - 2x} \tag{7.16}$$

whose power series expansion is $G(x) = \sum_{n\geq 2} 2^{n-2} x^n$.

□

Proof of Theorem 7.11. It is immediate from Proposition 7.12. □

Remark 7.17. Let T_n denote the Frobenius character $\mathrm{ch}(\chi^{n-1,n})$, and let $Y = Y(x)$ defined by

$$Y = \sum_{n=0}^{\infty} T_n \frac{x^n}{n!} \qquad (T_0 = 1)$$

denote its generating function. We think of generating series $H = \sum_{n=0}^{\infty} s_{(n)}$ as an operator acting (on the left) on series of symmetric functions by plethysm. By (7.1) we see that the Frobenius character series of $\chi^{k,n}$ is $\frac{1}{k!} s_{(1)}(H \circ Y)^k$. Therefore, a lá Polya, the functional equation that is satisfied by Y is found to be

$$Y = s_{(1)}x + s_{(1)}xH \circ Y + \frac{s_{(1)}x(H \circ Y)^2}{2!} + \frac{s_{(1)}x(H \circ Y)^3}{3!} + \cdots = xs_{(1)}e^{H \circ Y}.$$

Equivalently, we have

$$\frac{Y(x)}{\mathcal{H}(Y(x))} = x, \tag{7.18}$$

where $\mathcal{H}$ is the operation of applying the operator H (plethysitically) on the left, exponentiating the result, and then multiplying it by $s_{(1)}$. Thus, the solution to the Lagrange inversion problem (7.18) gives us the Frobenius character of $\chi^{n-1,n}$.

8 Non-attacking nilpotent rooks

The non-attacking rook placements are special file placements, hence they correspond to special labeled forests. Indeed, the forest of a non-attacking nilpotent k-rook placement has exactly $n - k$ connected components, each of which is a tree whose vertices have at most one child. In other words, each connected component is a chain. Since the ordering of these components does not play a role as far as the underlying labeled structure is concerned, the odun of the forest is completely determined by the sizes of the corresponding chains. Therefore, we have a one-to-one correspondence between oduns of nilpotent non-attacking k-rook placements on an $n \times n$ board and the partitions of n with exactly $n - k$ parts.

Let us denote the Frobenius character of the nilpotent non-attacking $n - k$-rook placements by $Z_{k,n}$. By the discussion above and by (7.1) we re-express $Z_{k,n}$ as in

$$Z_{k,n} = \sum_{(\lambda_1,\dots,\lambda_k)=(1^{m_1}2^{m_2}\dots)\vdash n} s_{(m_i)} \circ s_{(1)}^i, \tag{8.1}$$

where m_i is the number of times the part i occurs in λ. Let us denote by $\mathcal{R}_{k,n}$ the set of all non-attacking k-rook placements on $[n] \times [n]$. Since $\mathcal{R}_n = \bigcup_{k=1}^{n} \mathcal{R}_{k,n}$, the corresponding Frobenius character of nilpotent rook placements is $Z_n := \sum_{k=1} Z_{k,n}$.

Theorem 8.2. *The number of nilpotent non-attacking $n-k$-rook placements is given by*

$$\langle Z_{k,n}, s^n_{(1)}\rangle = \langle s^n_{(1)}, \sum_{(\lambda_1,\dots,\lambda_k)=(1^{m_1}2^{m_2}\dots)\vdash n} s_{(m_i)} \circ s^i_{(1)}\rangle.$$

Proof. If F_V denotes the Frobenius character of an S_n-module V, then $\dim V = \langle \mathcal{F}_V, s^n_{(1)}\rangle$. □

Let $\mathrm{Nil}(\mathcal{R}_{n-k,n})$ denote S_n-representation that is defined on the rank $n-k$ elements of $\mathrm{Nil}(\mathcal{R}_n)$.

Theorem 8.3. *Let n and k be two positive integers such that $n \geq 2$ and $1 \leq k \leq n-1$. In this case, the total number of irreducible representations in $Nil(\mathcal{R}_{n-k,n})$ is equal to*

$$p(k,n) = \#\textit{number of partitions of } n \textit{ with } k \textit{ non-zero parts.}$$

Similarly, the multiplicity of the sign representation in $Nil(\mathcal{R}_{n-k,n})$ is equal to the number of partitions of n with k parts $(\lambda_1,\dots,\lambda_k)$ such that for each $i \in [k]$ either λ_i is even, or $\lambda_j \neq \lambda_i$ for all $j \in [n]-\{i\}$.

Proof. The proof follows from the discussion above and Proposition 7.10. □

9 Final remarks

By Theorem 8.3, if τ is the labeled forest of a non-attacking $n-k$ rook placement, then the corresponding representation $o(\tau)$ has a well-defined partition type. Let $\lambda = (\lambda_1,\dots,\lambda_k) = (1^{m_1}2^{m_2}\dots) \vdash n$ denote the partition type of the odun $o(\tau)$. It is desirable to understand the dimension of $o(\tau)$ in terms of λ. Since we are working with permutation representations, the cardinality of an S_n-set gives the dimension of the corresponding representation, therefore, the number of labelings of the underlying forest for τ is equal to the dimension odun $o(\tau)$. After these remarks it is not difficult to see that the dimension of $o(\tau)$ is given by the formula

$$\langle F_{o(\tau)}, s^n_{(1)}\rangle = \frac{n!}{m_1!m_2!\cdots}. \tag{9.1}$$

A rooted tree σ can be viewed as a (graded) poset with a unique maximal element. If $a \in \sigma$ is a vertex, then its *hook* is defined to be

$$\sigma_a := \{b \in \sigma :\ b \leq a\}. \tag{9.2}$$

The corresponding *hook-length* is defined as the cardinality $h(a) := |\sigma_a|$. A *natural labeling* on σ is a bijection $g : o \to [n]$ such that $a < b$ implies $g(a) > g(b)$. The famous 'hook-length formula' of Knuth [17] which is proven by Sagan in his thesis [24] asserts that the number of natural labelings of σ is equal to

$$f^\sigma = \frac{n!}{\prod_{a\in\sigma} h(a)}. \tag{9.3}$$

Let τ be as in the previous paragraph so that its odun $o(\tau)$ consists of m_1 chains of length 1, m_2 chains of length 2, and so on. We add a new root to τ to obtain a rooted tree $\sigma = \sigma_\tau$. Since $\langle F_{o(\sigma)}, s^{n+1}_{(1)} \rangle = \langle F_\tau, s^n_{(1)} \rangle$, by (9.1) we know the dimension of $o(\sigma)$, which is $\dim F_{o(\sigma)} = \frac{n!}{m_1! m_2! \cdots}$. On the other hand, by (9.3), we see that $f^\sigma = \frac{(n+1)!}{(2!)^{m_1}(3!)^{m_2}\cdots}$, which is different than $\dim F_o$. In the next subsection we explain a more general dimension formula for the dimension of a forest representation.

9.1 Dimension of a forest representation

Let σ be a labeled forest thought of as a poset (with multiple maximal elements). If a is a vertex of σ, then we denote by σ^0_a the (rooted) subforest $\{b \in \sigma : \ b < a\}$. Recall that the notation σ_a stands for the hook (9.2) of a, which is in fact a subtree (of a subtree) of σ. Finally, let γ_a denote the set $\{a_1, \dots, a_r\}$, the complete list of children of a whose corresponding subtrees σ_{a_i} are mutually distinct, that is to say $\sigma_{a_i} \neq \sigma_{a_j}$ if $1 \leq i \neq j \leq r$. In this case, for $i = 1, \dots, r$ we denote by $m(a; a_i)$ the multiplicity of σ_{a_i} in σ^0_a.

The final result of our paper is the following formula for the dimension of the odun of σ.

Theorem 9.4. *If σ is either a tree on $n+1$ vertices or a forest on n vertices, then the dimension of the corresponding representation is*

$$\dim o(\sigma) = \frac{n!}{\prod_{a \in \sigma} \prod_{b \in \gamma_a} m(a; b)!}. \tag{9.5}$$

Proof. Note that if σ is a rooted tree on $n+1$ vertices and σ^0 denotes the forest obtained from σ by removing the root, then $\langle F_\sigma, s^{n+1}_{(1)} \rangle = \langle F_{\sigma^0}, s^n_{(1)} \rangle$. Therefore, it suffices to prove our claim for labeled trees only.

Let a be a vertex in σ and σ_b be a subtree where $b \in \gamma_a$. Then the permutation of the $m(a; b)$ copies of the σ_b does not change the labels of σ. See Figure 9.1 for a simple example. Therefore, the cardinality of the stabilizer subgroup of σ is given by $\prod_{a \in \sigma} \prod_{b \in \gamma_a} m(a; b)!$,

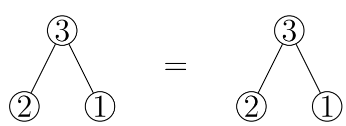

Figure 9.1: Permuting identical neighboring subtrees does not change the labeling.

hence the proof follows. $\square$

10 Decomposition tables

10.1 $n = 4$

$$\begin{aligned}
\mathcal{C}_{0,4} &= V_{(4)}\\
\mathcal{C}_{1,4} &= V_{(4)} \oplus V^2_{(3,1)} \oplus V_{(2,2)} \oplus V_{(2,1,1)}\\
\mathcal{C}_{2,4} &= V^3_{(4)} \oplus V^6_{(3,1)} \oplus V^5_{(2,2)} \oplus V^5_{(2,1,1)} \oplus V^2_{(1,1,1,1)}\\
\mathcal{C}_{3,4} &= V^4_{(4)} \oplus V^9_{(3,1)} \oplus V^5_{(2,2)} \oplus V^7_{(2,1,1)} \oplus V^2_{(1,1,1,1)}
\end{aligned}$$

10.2 $n = 5$

$$\begin{aligned}
\mathcal{C}_{0,5} &= V_{(5)}\\
\mathcal{C}_{1,5} &= V_{(5)} \oplus V^2_{(4,1)} \oplus V_{(3,2)} \oplus V_{(3,1,1)}\\
\mathcal{C}_{2,5} &= V^3_{(5)} \oplus V^7_{(4,1)} \oplus V^8_{(3,2)} \oplus V^6_{(3,1,1)} \oplus V^6_{(2,2,1)} \oplus V^3_{(2,1,1,1)} \oplus V_{(1,1,1,1,1)}\\
\mathcal{C}_{3,5} &= V^6_{(5)} \oplus V^{20}_{(4,1)} \oplus V^{22}_{(3,2)} \oplus V^{25}_{(3,1,1)} \oplus V^{19}_{(2,2,1)} \oplus V^{14}_{(2,1,1,1)} \oplus V^3_{(1,1,1,1,1)}\\
\mathcal{C}_{4,5} &= V^9_{(5)} \oplus V^{26}_{(4,1)} \oplus V^{28}_{(3,2)} \oplus V^{30}_{(3,1,1)} \oplus V^{24}_{(2,2,1)} \oplus V^{17}_{(2,1,1,1)} \oplus V^4_{(1,1,1,1,1)}
\end{aligned}$$

10.3 $n = 6$

$$\begin{aligned}
\mathcal{C}_{0,6} &= V_{(6)}\\
\mathcal{C}_{1,6} &= V_{(6)} \oplus V^2_{(5,1)} \oplus V_{(4,2)} \oplus V_{(4,1^2)}\\
\mathcal{C}_{2,6} &= V^3_{(6)} \oplus V^7_{(5,1)} \oplus V^9_{(4,2)} \oplus V^6_{(4,1^2)} \oplus V^3_{(3^2)} \oplus V^7_{(3,2,1)} \oplus V^3_{(3,1^3)} \oplus V^2_{(2^3)} \oplus V_{(2^2,1^2)} \oplus V_{(2,1^4)}\\
\mathcal{C}_{3,6} &= V^7_{(6)} \oplus V^{23}_{(5,1)} \oplus V^{35}_{(4,2)} \oplus V^{33}_{(4,1^2)} \oplus V^{19}_{(3^2)} \oplus V^{47}_{(3,2,1)} \oplus V^{24}_{(3,1^3)} \oplus V^{14}_{(2^3)} \oplus V^{21}_{(2^2,1^2)} \oplus V^9_{(2,1^4)} \oplus V^2_{(1^6)}\\
\mathcal{C}_{4,6} &= V^{16}_{(6)} \oplus V^{59}_{(5,1)} \oplus V^{96}_{(4,2)} \oplus V^{96}_{(4,1^2)} \oplus V^{46}_{(3^2)} \oplus V^{142}_{(3,2,1)} \oplus V^{83}_{(3,1^3)} \oplus V^{43}_{(2^3)} \oplus V^{68}_{(2^2,1^2)} \oplus V^{36}_{(2,1^4)} \oplus V^6_{(1^6)}\\
\mathcal{C}_{5,6} &= V^{20}_{(6)} \oplus V^{75}_{(5,1)} \oplus V^{114}_{(4,2)} \oplus V^{117}_{(4,1^2)} \oplus V^{59}_{(3^2)} \oplus V^{170}_{(3,2,1)} \oplus V^{96}_{(3,1^3)} \oplus V^{49}_{(2^3)} \oplus V^{83}_{(2^2,1^2)} \oplus V^{42}_{(2,1^4)} \oplus V^8_{(1^6)}
\end{aligned}$$

References

[1] Kürşat Aker and Mahir Bilen Can. From parking functions to Gelfand pairs. *Proc. Amer. Math. Soc.*, 140(4):1113–1124, 2012.

[2] Fred Butler. Rook theory and cycle-counting permutation statistics. *Adv. in Appl. Math.*, 33(4):655–675, 2004.

[3] Fred Butler, Mahir Bilen Can, James Haglund, and Jeff Remmel. *Rook Theory.* http://www.math.ucsd.edu/ remmel/files/Book.pdf, 2015. In preparation.

[4] Louis Comtet. *Advanced combinatorics*. D. Reidel Publishing Co., Dordrecht, enlarged edition, 1974. The art of finite and infinite expansions.

[5] William F. Doran, IV. A plethysm formula for $p_\mu(\underline{x}) \circ h_\lambda(\underline{x})$. *Electron. J. Combin.*, 4(1):Research Paper 14, 10 pp. (electronic), 1997.

[6] Olexandr Ganyushkin and Volodymyr Mazorchuk. *Classical finite transformation semigroups*, volume 9 of *Algebra and Applications*. Springer-Verlag London, Ltd., London, 2009. An introduction.

[7] A. M. Garsia and M. Haiman. A remarkable q,t-Catalan sequence and q-Lagrange inversion. *J. Algebraic Combin.*, 5(3):191–244, 1996.

[8] A. M. Garsia and J. B. Remmel. Q-counting rook configurations and a formula of Frobenius. *J. Combin. Theory Ser. A*, 41(2):246–275, 1986.

[9] Ira M. Gessel. Generalized rook polynomials and orthogonal polynomials. In *q-series and partitions (Minneapolis, MN, 1988)*, volume 18 of *IMA Vol. Math. Appl.*, pages 159–176. Springer, New York, 1989.

[10] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete mathematics*. Addison-Wesley Publishing Company, Reading, MA, second edition, 1994. A foundation for computer science.

[11] J. Haglund, M. Haiman, N. Loehr, J. B. Remmel, and A. Ulyanov. A combinatorial formula for the character of the diagonal coinvariants. *Duke Math. J.*, 126(2):195–232, 2005.

[12] James Haglund. Rook theory and hypergeometric series. *Adv. in Appl. Math.*, 17(4):408–459, 1996.

[13] James Haglund. *The q,t-Catalan numbers and the space of diagonal harmonics*, volume 41 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2008. With an appendix on the combinatorics of Macdonald polynomials.

[14] Mark Haiman. Combinatorics, symmetric functions, and Hilbert schemes. In *Current developments in mathematics, 2002*, pages 39–111. Int. Press, Somerville, MA, 2003.

[15] Roger Howe. $(\mathrm{GL}_n, \mathrm{GL}_m)$-duality and symmetric plethysm. *Proc. Indian Acad. Sci. Math. Sci.*, 97(1-3):85–109 (1988), 1987.

[16] Irving Kaplansky and John Riordan. The problem of the rooks and its applications. *Duke Math. J.*, 13:259–268, 1946.

[17] Donald E. Knuth. *The art of computer programming. Vol. 3*. Addison-Wesley, Reading, MA, 1998. Sorting and searching, Second edition [of MR0445948].

[18] A. Laradji and A. Umar. On the number of nilpotents in the partial symmetric semigroup. *Comm. Algebra*, 32(8):3017–3023, 2004.

[19] D. E. Littlewood. Invariant theory, tensors and group characters. *Philos. Trans. Roy. Soc. London. Ser. A.*, 239:305–365, 1944.

[20] Nicholas A. Loehr and Jeffrey B. Remmel. A computational and combinatorial exposé of plethystic calculus. *J. Algebraic Combin.*, 33(2):163–198, 2011.

[21] I. G. Macdonald. *Symmetric functions and Hall polynomials*. Oxford Mathematical Monographs. The Clarendon Press, Oxford University Press, New York, second edition, 1995. With contributions by A. Zelevinsky, Oxford Science Publications.

[22] Mohan S. Putcha. *Linear algebraic monoids*, volume 133 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1988.

[23] Lex E. Renner. *Linear algebraic monoids*, volume 134 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin, 2005. Invariant Theory and Algebraic Transformation Groups, V.

[24] Bruce Eli Sagan. *PARTIALLY ORDERED SETS WITH HOOKLENGTHS - AN ALGORITHMIC APPROACH*. ProQuest LLC, Ann Arbor, MI, 1979. Thesis (Ph.D.)–Massachusetts Institute of Technology.

[25] Richard P. Stanley and John R. Stembridge. On immanants of Jacobi-Trudi matrices and permutations with restricted position. *J. Combin. Theory Ser. A*, 62(2):261–279, 1993.