The Ubiquitous Young Tableau

Bruce E. Sagan
Department of Mathematics
Michigan State University
East Lansing, MI 48824-1027

August 5, 1988

Proposed running head:

UBIQUITOUS YOUNG TABLEAU

Send proofs to:

Bruce E. Sagan
Department of Mathematics
Michigan State University
East Lansing, MI 48824-1027

Abstract

Young tableaux have found extensive application in combinatorics [Vie 84], group representations [Jam 78], invariant theory [DRS 74, DKR 78], symmetric funcions [Mad 79], and the theory of algorithms [Knu 73, pages 48-72]. This paper is an expository treatment of some of the highlights of tableaux theory. These include the hook and determinantal formulae for enumeration of both standard and generalized tableaux, their connection with irreducible representations of matrix groups, and the Robinson-Schensted-Knuth algorithm.

1 Three families of tableaux

Young tableaux were first introduced in 1901 by the Reverend Alfred Young [You 01, page 133] as a tool for invariant theory. Subsequently, he showed that they can give information about representations of symmetric groups. Since then, tableaux have played an important rôle in many areas of mathematics from enumerative combinatorics to algebraic geometry. This paper is a survey of some of these applications.

In recent years the number of tableaux of various types has been increasing at an impressive rate. To limit this paper to a reasonable length, our discussion will be restricted to three fundamental families of tableaux: ordinary, shifted and oscillating.

The rest of this section will be devoted to the definitions and notation need to describe these arrays. In Section 2 we present the hook and determinantal formulae for enumeration of standard tableaux. The third section examines the connection with representations of the symmetric group. The Robinson-Schensted algorithm appears in Section 4 as a combinatorial way of explaining the decomposition of the regular representation. The next four sections rework the material from the first four using generalized tableaux (those with repeated entries), representations of general linear and symplectic groups, and the theory of symmetric functions. Section 9 is a brief exposition of some open problems.

1.1 Ordinary tableaux

In what follows, $\mathbf{N}$ and $\mathbf{P}$ stand for the non-negative and positive integers respectively. A *partition* λ *of* $n \in \mathbf{N}$, written $\lambda \vdash n$, is a sequence of positive integers $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_l)$ in weakly decreasing order such that $\sum_{i=1}^{l} \lambda_i = n$. The λ_i are called the *parts of* λ. The unique partition of 0 is $\lambda = \phi$. The *shape of* λ is an array of boxes (or dots or cells) with l left-justified rows and λ_i boxes in row i. We will use λ to represent both the partition and its shape, while (i, j) will denote the cell in row i and column j. By way of illustration, the following figure shows the shape of the partition $\lambda = (2, 2, 1) \vdash 5$ with cell (3,1) displayed as a diamond[1].

[1]In deference to Alfred Young's nationality, we have chosen to draw partition shapes in the English style, i.e., as if they were part of a matrix. The reader should be aware that some mathematicians (notably the French) prefer to use the conventions of coordinate geometry where λ_1 cells are placed along the x-axis, λ_2 cells are placed along the line $y = 1$, etc. To them and to René Descartes, we abjectly apologize.

$$\begin{array}{cc} \square & \square \\ \square & \square \\ \diamond & \end{array} \quad .$$

A *standard Young tableau (SYT) of shape* λ, denoted P, is obtained by filling the cells of $\lambda \vdash n$ with the integers from 1 to n so that

1. each integer is used exactly once, and
2. the rows and columns increase.

We let $p_{i,j}$ denote the elemenet of P in cell (i, j). There are 5 SYT of shape $(2, 2, 1)$:

$$\begin{array}{cc} 1 & 2 \\ 3 & 4 \\ 5 & \end{array} \; , \; \begin{array}{cc} 1 & 2 \\ 3 & 5 \\ 4 & \end{array} \; , \; \begin{array}{cc} 1 & 3 \\ 2 & 4 \\ 5 & \end{array} \; , \; \begin{array}{cc} 1 & 3 \\ 2 & 5 \\ 4 & \end{array} \; , \; \begin{array}{cc} 1 & 4 \\ 2 & 5 \\ 3 & \end{array}$$

and the first tableau has $p_{3,1} = 5$. Letting f^{λ} be the number of SYT of shape λ, we see that $f^{(2,2,1)} = 5$.

1.2 Shifted tableaux

A partition $\lambda^* = (\lambda_1, \lambda_2, \cdots, \lambda_l)$ of n is *strict*, $\lambda^* \models n$, if $\lambda_1 > \lambda_2 > \cdots > \lambda_l$. The *shifted shape of* λ^* is like the ordinary shape except that row i starts with its leftmost box in position (i, i). The fact that the parts of λ^* strictly decrease assures that shifting the rows in this manner does not cause any cells to stick out from the right-hand boundary. As an example, the shifted shape of $\lambda^* = (4, 2, 1) \models 7$ is:

$$\begin{array}{cccc} \square & \square & \square & \square \\ & \square & \square & \\ & & \square & \end{array} \quad .$$

A *standard shifted Young tableau (SST) of shape* λ^*, P^*, is a filling of the shifted shape λ^* satisfying the same two conditions as for an SYT. The notation $p^*_{i,j}$ should be self-explanatory. The number of SST of shape λ^* is denoted g^{λ}. The list below displays all SST of shifted shape $(4, 2, 1)$, demonstrating that $g^{(4,2,1)} = 7$.

$$\begin{array}{cccc} 1 & 2 & 3 & 4 \\ & 5 & 6 & \\ & & 7 & \end{array} \; , \; \begin{array}{cccc} 1 & 2 & 3 & 5 \\ & 4 & 6 & \\ & & 7 & \end{array} \; , \; \begin{array}{cccc} 1 & 2 & 3 & 6 \\ & 4 & 5 & \\ & & 7 & \end{array} \; , \; \begin{array}{cccc} 1 & 2 & 3 & 7 \\ & 4 & 5 & \\ & & 6 & \end{array} \; ,$$

$$\begin{array}{cccc} 1 & 2 & 4 & 5 \\ & 3 & 6 & \\ & & 7 & \end{array} \; , \; \begin{array}{cccc} 1 & 2 & 4 & 6 \\ & 3 & 5 & \\ & & 7 & \end{array} \; , \; \begin{array}{cccc} 1 & 2 & 4 & 7 \\ & 3 & 5 & \\ & & 6 & \end{array} \; .$$

1.3 Oscillating tableaux

To motivate the definition of an oscillating tableau, we must first look again at the definition of an SYT. An SYT P of shape λ gives rise to a sequence of shapes $\phi = \lambda^0 \subset \lambda^1 \subset \lambda^2 \subset \cdots \subset \lambda^n = \lambda$ where λ^m is the shape containing the numbers $1, 2, \cdots, m$ in P, i. e. , λ^m is obtained from λ^{m-1} by adding the cell (i, j) such that $p_{i,j} = m$. For example if

$$P = \begin{array}{ll} 1 & 3 \\ 2 & 5 \\ 4 & \end{array}$$

then the corresponding sequence is

$$\phi \subset \square \subset \begin{array}{l} \square \\ \square \end{array} \subset \begin{array}{ll} \square & \square \\ \square & \end{array} \subset \begin{array}{ll} \square & \square \\ \square & \\ \square & \end{array} \subset \begin{array}{ll} \square & \square \\ \square & \square \\ \square & \end{array} = \lambda.$$

Conversely it should be clear that any sequence of shapes starting with ϕ and adding a box at each stage corresponds to a SYT. The definition of oscillating tableau generalizes this concept.

An *oscillating Young tableau (OYT) of shape* λ *and length* k, $\tilde{P}_k^\lambda$, is a sequence of shapes $(\phi = \lambda^0, \lambda^1, \cdots, \lambda^k = \lambda)$ such that λ^m is obtained from λ^{m-1} by adding or subtracting a cell. Oscillating tableaux are also called *up-down* or *alternating* tableaux. We let $\tilde{f}_k^\lambda$ denote the number of OYT of shape λ and length k. If $\lambda = (1)$ and $k = 3$ then a complete list of the corresponding OYT is:

$$(\phi\,,\ \square\,,\ \phi\,,\ \square\,)\ ;\ (\phi\,,\ \square\,,\ \square\,\square\,,\ \square\,)\ ;\ (\phi\,,\ \square\,,\ \begin{array}{l} \square \\ \square \end{array},\ \square\,)\ .$$

It follows that $\tilde{f}_3^{(1)} = 3$.

2 Enumeration of tableaux

It would be useful to have an expression for the number of tableaux of a given shape, since enumeration by hand (as in the previous section) rapidly becomes unwieldy as n increases. There are two principle formulae of this type, one involving products (the hook formula) and one involving determinants.

2.1 Ordinary tableaux

If (i, j) is a cell in the shape of λ then it has *hook*

$$H_{i,j} = \{(i,j)\} \ \cup \ \{(i,j') \mid j' > j\} \ \cup \ \{(i',j) \mid i' > i\}$$

with corresponding *hooklength* $h_{i,j} = |H_{i,j}|$ (where $|\cdot|$ stands for cardinality). The sets $\{(i,j')|j' > j\}$ and $\{(i',j)|i' > i\}$ are called the *arm* and *leg* of the hook respectively. If $\lambda = (6, 5, 3, 2)$ then the diamonds in

$$\begin{array}{cccccc} \square & \square & \diamond & \diamond & \diamond & \diamond \\ \square & \square & \diamond & \square & \square & \\ \square & \square & \diamond & & & \\ \square & \square & & & & \end{array}$$

represent the hook $H_{1,3}$ with hooklength $h_{1,3} = 6$. The famous *hook formula* expresses the number of SYT in terms of hooklengths.

Theorem 2.1.1 (Frame-Robinson-Thrall [FRT 54]) *If $\lambda \vdash n$ then*

$$f^\lambda = \frac{n!}{\prod_{(i,j)\in\lambda} h_{i,j}}.$$

Before discussing various proofs of this theorem, let us look at an example to see how easy it is to apply the hook formula. If $\lambda = (2, 2, 1) \vdash 5$ then the hooklengths of λ are given in the diagram

$$\begin{array}{cc} 4 & 2 \\ 3 & 1 \\ 1 & \end{array}$$

where $h_{i,j}$ is placed in cell (i, j). Thus $f^{(2,2,1)} = 5!/4 \cdot 3 \cdot 2 \cdot 1^2 = 5$ which agrees with our previous computation in Section 1.1. There are many different proofs of the hook formula; we outline a few of them next.

Proof sketches.

1. (inductive) It is easy to prove the hook formula by induction on n. Unfortunately this type of proof gives no inkling of *why* the hooklengths should play a role.

2. (probabilistic) Greene, Nijenhuis, and Wilf [GNW 79] have given a beautiful probabilistic proof where the hooks do enter in a very strong way. The general idea is this. Fix a shape $\lambda \vdash n$. If we can find an algorithm that produces any SYT P with probability $prob(P) = \prod h_{i,j}/n!$, then we will be done since the distribution is uniform. In what follows a *corner cell* is $(i,j) \in \lambda$ such that $(i+1,j),(i,j+1) \notin \lambda$.

   (a) Pick $(i,j) \in \lambda$ with probability $1/n$.

   (b) While (i,j) is not a corner cell do **begin**

      i. pick a cell $c \in H_{i,j} - \{(i,j)\}$ with probability $1/(h_{i,j}-1)$;

      ii. $(i,j) := c$, i.e., c becomes the new value for (i,j) **end**.

   (c) Give the label n to the corner cell (i',j') that you have reached.

   (d) Go back to step (a) with $\lambda := \lambda - \{(i',j')\}$ and $n := n-1$. Repeat this outer loop until all cells of λ are labeled.

   It should be clear that this procedure gives a standard labeling of λ . It is less obvious (though not hard to prove) that all labelings are equally likely and of the right probability. The interesed reader can consult [GNW 79] for the details.

3. (combinatorial) Franzblau and Zeilberger [F-Z 82] were the first to come up with a combinatorial proof of the hook formula. Rewriting the equation as $n! = f^\lambda \cdot \prod h_{i,j}$, we see that it suffices to find a bijection $S \longleftrightarrow (P,H)$ where S is an arbitrary filling of λ with $1,2,\cdots,n$ (rows and columns need not increase), P is a SYT of shape λ , and H is a *pointer tableau* of shape λ , i.e., a placement of (computer science-type) pointers in λ such that the pointer in cell (i,j) points to some cell of $H_{i,j}$. Roughly, given a scrambled tableau S we wish to rearrange its entries to form a SYT P with the pointer tableau H keeping track of the unscrambling process. While this idea is simple, the actual bijection is long and difficult. Subsequently Zeilberger [Zei 84] found a way to turn the Greene-Nijenhuis-Wilf proof into a bijection, but the details are still not as pleasant as one would like. $\square$

It is unfortunate that a simple combinatorial statement like the hook formula has no simple combinatorial proof. (It seems as if all the simple proofs are not combinatorial and all the combinatorial proofs are not simple!) To find such a proof is one of the tantalizing open problems in this area.

While the hook formula is relatively recent, the determinantal formula goes back to Frobenius and Young. In what follows, $1/r! = 0$ if $r < 0$.

Theorem 2.1.2 *If* $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_l) \vdash n$ *then*

$$f^\lambda = n! \cdot det[1/(\lambda_i - i + j)!]$$

where the determinant is $l \times l$*.*

The simplest way to remember the denominators in the determinant is to note that the parts of λ are found along the main diagonal. The other entries in a given row are computed by increasing or decreasing the number (inside the factorial) by 1 for each step taken to the right or left respectively. If we apply this formula to our running example where $\lambda = (2, 2, 1)$ we find

$$5! \cdot \begin{vmatrix} 1/2! & 1/3! & 1/4! \\ 1/1! & 1/2! & 1/3! \\ 0 & 1/0! & 1/1! \end{vmatrix} = 5$$

which has not changed since our last computation.

Proof sketches (of the determinantal formula).

1. (inductive) If $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_l)$ then it is easy to see directly from the definitions that $\lambda_i - i = h_{i,1} - l$. Hence it is enough to show that

   $$f^\lambda = n! \cdot det[1/(h_{i,1} - l + j)!].$$

   This can be done using induction and the hook formula.

2. (combinatorial) A SYT can be represented as a family of non-intersecting lattice paths in the plane. Such families are also counted by determinants, as shown by Gessel [unpublished manuscript]. Remmel [Rem 82] combined these ideas with the Garsia-Milne involution principle [G-M 81] to give a bijective proof of both the hook and determinantal formulae. Gessel and Viennot [G-V 85, G-V ip] have extended this idea to a multitude of interesting applications. $\square$

2.2 Shifted tableaux

A shifted hook is like a hook with an extra appendage. The *shifted hook* of $(i,j) \in \lambda^*$ is

$$H_{i,j}{}^* = \{(i,j)\} \;\cup\; \{(i,j') \mid j' > j\} \;\cup\; \{(i',j) \mid i' > i\} \;\cup\; \{(j+1,j') | j' > j\}$$

with *hooklength* $h^*_{i,j} = |H_{i,j}{}^*|$. The diamonds below outline the hook $H^*_{2,3}$ of the partition $\lambda^* = (6,5,3,2,1)$ which has hooklength $h^*_{2,3} = 7$:

$$\begin{array}{cccccc}
\square & \square & \square & \square & \square & \square \\
 & \square & \diamond & \diamond & \diamond & \diamond \\
 & & \diamond & \square & \square & \\
 & & & \diamond & \diamond & \\
 & & & & \square &
\end{array} \quad .$$

One way to motivate this definition is to note that if λ^* is a shifted shape then one can paste together λ^* and $\lambda^{*\,t}$ (where t denotes the transpose) to form a left justified shape λ. If $(i,j) \in \lambda^*$ then $H_{i,j}{}^* \subseteq \lambda^*$ is the same as $H_{i,j} \subseteq \lambda$ except that the bottom part of its leg has been twisted. For example, if we use $\lambda^* = (6,5,3,2,1)$ as before and represent $\lambda^{*\,t}$ using circles, then the shifted hook above corresponds to the normal hook in

$$\begin{array}{ccccccc}
\circ & \square & \square & \square & \square & \square & \square \\
\circ & \circ & \square & \diamond & \diamond & \diamond & \diamond \\
\circ & \circ & \circ & \diamond & \square & \square & \\
\circ & \circ & \circ & \diamond & \square & \square & \\
\circ & \circ & \circ & \diamond & \circ & \square & \\
\circ & \circ & & & & &
\end{array} \quad .$$

The shifted analog of the hook formula is

Theorem 2.2.1 ([Thr 52]) *If* $\lambda^* \models n$ *then*

$$g^\lambda = \frac{n!}{\prod_{(i,j)\in\lambda^*} h^*_{i,j}}. \quad \square$$

As an example, $\lambda^* = (4,2,1) \models 7$ has shifted hooklengths

$$\begin{array}{cccc}
6 & 5 & 4 & 1 \\
 & 3 & 2 & \\
 & & 1 &
\end{array}$$

and so $g^{(4,2,1)} = 7!/6 \cdot 5 \cdot 4 \cdot 3 \cdot 2 \cdot 1^2 = 7$ as noted in Section 1.2.

It is easy to give an inductive proof of this result. A probabilistic analog of the [GNW 79] proof was given by Sagan [Sag 80]. Surprisingly, although the algorithm for producing a tableau at random is identical (merely use shifted hooks in place of normal hooks), the proof that every tableau of a given shape is equally likely is much more complicated. It would be nice to find a simple proof that the shifted algorithm works. It is also an open problem to find analogs of the combinatorial proofs of [F-Z 82] and [Zei 84].

The reader is probably wondering why Thrall's [Thr 52] paper with the shifted hook formula appeared two years earlier than his article with Frame and Robinson [FRT 54] containing the unshifted version. In fact the 1952 paper contains an expression for g^λ which is mid-way between the shifted versions of the hook and determinantal formulae (although hooks are never mentioned explicitly) and from which either can be derived by simple manipulations. This brings us to the determinantal version.

Theorem 2.2.2 ([Thr 52]) *If* $\lambda^* = (\lambda_1, \lambda_2, \cdots, \lambda_l) \models n$ *then*

$$g^\lambda = \frac{n!}{\prod_{i<j}(\lambda_i + \lambda_j)} \cdot det[1/(\lambda_i - l + j)!]$$

where the determinant is $l \times l$. $\square$

Note that the parts of λ^* are now found in the last column of the determinant. Also the extra product in the denominator is precisely the set of hooklengths for the cells $(i, j) \in \lambda^*$ such that $i < j$. Thrall proved this theorem by induction. It seems probable that the techniques of Remmel [Rem 82] and Gessel-Viennot [G-V 85, G-V ip] can also be used.

2.3 Oscillating tableaux

If $\lambda \vdash n$ and $\tilde{f}_k^\lambda \neq 0$ then we must take at least k steps to reach λ and so $k \geq n$. Furthermore, the number of extra additions and subtractions of cells must cancel out, so $k \equiv n \pmod 2$ and thus $k - n = 2d$ for some $d \in \mathbf{N}$. With these preliminaries we can state a formula of Sundaram for the number of OYT of given length and shape.

Theorem 2.3.1 ([Sun 86]) *If* $\lambda \vdash n$ *and* $k - n = 2d$ *for some* $d \in \mathbf{N}$*, then*

$$\tilde{f}_k^\lambda = \binom{k}{n}(2d)!!\ f^\lambda$$

where $\binom{k}{n}$ *is a binomial coefficient and* $(2d)!! = 1 \cdot 3 \cdot 5 \cdot 7 \cdots (2d-1)$ *, i.e.,* $(2d)!!$ *is the number of fixedpoint-free involutions in the symmetric group* S_{2d}.

Proof sketch. It suffices to find a bijection

$$\tilde{P} \longleftrightarrow (P, \pi)$$

where $\tilde{P}$ is an OYT of shape λ and length k, P is an 'SYT' of shape λ , and π is a fixedpoint-free involution in 'S_{2d}'. The reason for the quotes is that the tableau need not contain the integers from 1 to n and the involution need not be a permutation of 1 to $2d$. Rather, the integers from 1 to k are to be partitioned into 2 subsets of size n and $k - n = 2d$ (accounting for the binomial coefficient), with the elements of P taken from the first subset and those of π from the second.

The general idea is that P keeps track of those steps in the construction of $\tilde{P}$ where a box of λ is added for the last time, while π stores information about deletions. See [Sun 86] for details. □

3 Representations of groups

Since one of Young's original applications for his tableaux came from group representation theory, it behoves us to look at the connection. In what follows, let G be a group and let V be a finite dimensional vector space over the complex numbers, $\mathbf{C}$.

3.1 Ordinary representations

A *representation of* G is a homomorphism $\rho : G \to GL(V)$ where $GL(V)$ is the general linear group of V, i.e., the group of all invertible linear transformations from V to itself. Alternatively, a representation may be viewed as a vector space V together with an action of G on V by invertible linear transformations. The space V is called a *G-module* and if $g \in G$, $v \in V$ then the action $\rho(g)v$ is abbreviated to $g \cdot v$ or just gv. We call dim V the *degree* of the representation.

As an example, consider any group G and let $V = \mathbf{C}$. Then the map that sends every $g \in G$ to the identity linear transformation (i.e., $gv = v$ for all $g \in G$ and $v \in V$) is a representation called the *trivial* representation. The trivial representation has degree 1.

For a more substantive example, let G be the symmetric group S_n and let V be the set of all formal linear combinations

$$V = \{c_1\vec{1} + c_2\vec{2} + \cdots + c_n\vec{n} \mid c_k \in \mathbf{C} \text{ for all } k\}$$

which is a vector space over $\mathbf{C}$ with basis $\mathcal{B} = \{\vec{1}, \vec{2}, \cdots, \vec{n}\}$. If $\pi \in S_n$ then we define the action of π on a basis vector by letting

$$\pi(\vec{k}) = \overrightarrow{\pi(k)} .$$

Thus the matrix of π in the canonical basis is just the usual permutation matrix associated with π , e.g., if $\pi = (1,2)(3) \in S_3$ then $\pi(\vec{1}) = \vec{2}$, $\pi(\vec{2}) = \vec{1}$ and $\pi(\vec{3}) = \vec{3}$ so that

$$\rho(\pi) = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

in the basis $\mathcal{B}$. This representation is called the *natural* or *defining* representation of S_n and is of degree n.

A G-module V is called *irreducible* if there is no proper subspace W of V which is invariant under the set of linear transformations $\rho(G) = \{\rho(g) \mid g \in G\}$. Equivalently, V is irreducible if it has no basis $\mathcal{B}$ that simultaneously brings all the matrices to block form:

$$\rho(g) = \begin{bmatrix} A_g & B_g \\ 0 & C_g \end{bmatrix}$$

for all $g \in G$.

Obviously the trivial representation is irreducible, as are all degree one representations. On the other hand, the natural representation of S_n is not irreducible for $n \geq 2$ since the one-dimensional subspace generated by $\vec{1} + \vec{2} + \cdots + \vec{n}$ is invariant.

The irreducible representations of a group are important because they are the building blocks of all other representations under certain conditions. A G-module is said to be *completely reducible* if it is a direct sum of irreducible G-modules. Mashke's theorem gives us a large supply of completely reducible modules.

Theorem 3.1.1 *If G is a finite group then every G-module V is completely reducible.* $\square$

Although we do not have room here to prove the results that we will need from representation theory, the reader is encouraged to consult the excellent text of Ledermann [Led 77] or the up-coming book of Sagan [Sag ip] .

The next question to ask is: given G, how many irreducible G-modules are there? First, however, we must know when two modules are the same. We say that G-modules V and W are *equivalent*, written $V \cong W$, if there is a vector space isomorphism $\phi : V \rightarrow W$ that preserves the action of G, i.e.,

$$\phi(gv) = g\phi(v) \text{ for all } v \in V, g \in G.$$

Theorem 3.1.2 *If G is finite then the number of inequivalent irreducible G-modules is equal to the number of conjugacy classes of G.* $\square$

If $G = S_n$ then a conjugacy class consists of all permutations of a given cycle-type. But a cycle-type is just a partition of n.

Corollary 3.1.3 *The number of inequivalent irreducible S_n-modules is the number of partitions of n.* $\square$

So to find the number of irreducible S_3-modules we merely list all partitions of 3:

$$(3) = \begin{matrix} \square & \square & \square \end{matrix}\ , \ (2,1) = \begin{matrix} \square & \square \\ \square & \end{matrix}\ , \ (1,1,1) = \begin{matrix} \square \\ \square \\ \square \end{matrix}\ .$$

Since there are 3 partitions, there are 3 irreducible modules for S_3.

The irreducible S_n-module indexed by $\lambda \vdash n$ is usually denoted S^λ and called the *Specht module corresponding to* λ . It would be nice to know their dimensions.

Theorem 3.1.4 *If $\lambda \vdash n$ then* $\dim S^\lambda = f^\lambda$. $\square$

Returning to S_3, we can compute the dimensions of each irreducible by listing all the SYT of the appropriate shape (or by using the hook or determinantal formulae):

$$\begin{aligned} (3) &: \begin{matrix} 1 & 2 & 3 \end{matrix}\ ; \\ \\ (2,1) &: \begin{matrix} 1 & 2 \\ 3 & \end{matrix}\ , \quad \begin{matrix} 1 & 3 \\ 2 & \end{matrix}\ ; \\ \\ (1,1,1) &: \begin{matrix} 1 \\ 2 \\ 3 \end{matrix}\ ; \end{aligned}$$

so $\dim S^{(3)} = 1$, $\dim S^{(2,1)} = 2$, and $\dim S^{(1,1,1)} = 1$. Now for any n we have $\dim S^{(n)} = \dim S^{(1^n)} = 1$ where $(1^n) = (\overbrace{1, \cdots, 1}^{n})$. It turns out that $S^{(n)}$ corresponds to the trivial representation and $S^{(1^n)}$ corresponds to the one-dimensional *sign* representation that sends every $\pi \in S_n$ to the matrix $[sgn(\pi)]$.

For any finite group $G = \{g_1, g_2, \cdots, g_m\}$, the *group algebra* is the vector space

$$\mathbf{C}(G) = \{c_1\vec{g}_1 + c_2\vec{g}_2 + \cdots + c_m\vec{g}_m \mid c_k \in \mathbf{C} \text{ for all } k\}.$$

Clearly $\mathbf{C}(G)$ is a G-module under the action $h \cdot \vec{g} = \overrightarrow{hg}$ for all $g, h \in G$. Hence we can ask how $\mathbf{C}(G)$ decomposes into irreducibles.

Theorem 3.1.5 *Let G be finite and let $S^1, S^2, \cdots, S^c$ be a complete list of inequivalent irreducible G-modules. Then*

$$\mathbf{C}(G) \cong \bigoplus_{i=1}^{c} m_i S^i$$

where $m_i = \dim S^i$, i.e., every irreducible module appears in $\mathbf{C}(G)$ with multiplicity equal to its dimension. $\square$

Taking dimensions on both sides of the previous equation we obtain:

Corollary 3.1.6

$$|G| = \sum_{i=1}^{c} m_i^2. \ \square$$

Finally, speciallizing to the symmetric group yields:

Corollary 3.1.7

$$n! = \sum_{\lambda \vdash n} (f^\lambda)^2. \ \square$$

3.2 Projective representations

A *projective representation* of a group G is a homomorphism $\rho : G \to PGL(V)$ where $PGL(V)$ is the projective general linear group, i.e., $GL(V)$ modulo the scalar multiples of the identity transformation. Below we will list the projective analogs of the results from the previous section. For a more complete discussion, see the articles of Stembridge [Ste 87] and Józefiak [Józ pr].

The irreducible projective representations of S_n are indexed by strict partitions λ^* . Unfortunately the indexing is no longer a 1-to-1 correspondence as in the ordinary case. Define the *length of* $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_l)$, denoted $l = l(\lambda)$, to be the number of parts of λ. This should not be confused with the length of an OYT. Now if $\lambda^* \models n$ has length l, then when $n - l$ is even there is a single irreducible projective S_n-module S_0^λ corresponding to λ^*, but when $n - l$ is odd there are two such: S_1^λ and S_{-1}^λ.

By way of illustration, consider S_6. The corresponding strict partitions, λ^*, are

$$\begin{array}{cccccc} \square & \square & \square & \square & \square & \square \end{array}, \quad \begin{array}{ccccc} \square & \square & \square & \square & \square \\ & \square & & & \end{array}, \quad \begin{array}{cccc} \square & \square & \square & \square \\ & \square & \square & \end{array}, \quad \begin{array}{ccc} \square & \square & \square \\ & \square & \square \\ & & \square \end{array}$$

with $n-l$ being $6-1=5$, $6-2=4$, $6-2=4$, and $6-3=3$ respectively. Thus the number of irreducible projective representations of S_6 is $2+1+1+2=6$. (It is an accident that in all the examples we have seen, the number of irreducible representations of the S_n in question is always n.)

As in the ordinary case, we can also compute the dimensions of the irreducibles using tableaux. In the following theorem, $\lfloor\cdot\rfloor$ is the greatest integer function (also called the floor or round-down function). The power of 2 enters because the Schur multiplier for the symmetric group has order 2.

Theorem 3.2.1 *If $\lambda^* \models n$ has length l then*

$$\dim S_i^\lambda = 2^{\lfloor\frac{n-l}{2}\rfloor} \cdot g^\lambda$$

where $i=0$ if $n-l$ is even and $i=\pm 1$ if $n-l$ is odd. $\square$

To finish our computation for S_6, we list the shifted tableaux with shapes given by the shapes above:

$$\begin{array}{rcl}
(6) & : & \begin{array}{l}1\ 2\ 3\ 4\ 5\ 6\ ;\end{array} \\[2ex]
(5,1) & : & \begin{array}{l}1\ 2\ 3\ 4\ 5\ , \\ \ \ 6\end{array} \quad \begin{array}{l}1\ 2\ 3\ 4\ 6\ , \\ \ \ 5\end{array} \\[2ex]
 & & \begin{array}{l}1\ 2\ 3\ 5\ 6\ , \\ \ \ 4\end{array} \quad \begin{array}{l}1\ 2\ 4\ 5\ 6\ ; \\ \ \ 3\end{array} \\[2ex]
(4,2) & : & \begin{array}{l}1\ 2\ 3\ 4\ , \\ \ \ 5\ 6\end{array} \quad \begin{array}{l}1\ 2\ 3\ 5\ , \\ \ \ 4\ 6\end{array} \quad \begin{array}{l}1\ 2\ 3\ 6\ , \\ \ \ 4\ 5\end{array} \\[2ex]
 & & \begin{array}{l}1\ 2\ 4\ 5\ , \\ \ \ 3\ 6\end{array} \quad \begin{array}{l}1\ 2\ 4\ 6\ ; \\ \ \ 3\ 5\end{array} \\[2ex]
(3,2,1) & : & \begin{array}{l}1\ 2\ 3\ , \\ \ \ 4\ 5 \\ \ \ \ \ 6\end{array} \quad \begin{array}{l}1\ 2\ 4\ . \\ \ \ 3\ 5 \\ \ \ \ \ 6\end{array}
\end{array}$$

Hence $\dim S_1^{(6)} = \dim S_{-1}^{(6)} = 2^{\lfloor 5/2\rfloor}\cdot 1 = 4$, $\dim S_0^{(5,1)} = 2^{\lfloor 4/2\rfloor}\cdot 4 = 16$, $\dim S_0^{(4,2)} = 2^{\lfloor 4/2\rfloor}\cdot 5 = 20$, and $\dim S_1^{(3,2,1)} = \dim S_{-1}^{(3,2,1)} = 2^{\lfloor 3/2\rfloor}\cdot 1 = 4$.

Because Corollary 3.1.6 continues to hold for projective representations we have

$$\begin{aligned} n! &= \sum_{\substack{\lambda^* \models n \\ n-l \text{ even}}} (\dim S_0^\lambda)^2 + \sum_{\substack{\lambda^* \models n \\ n-l \text{ odd}}} (\dim S_1^\lambda)^2 + (\dim S_{-1}^\lambda)^2 \\ &= \sum_{\substack{\lambda^* \models n \\ n-l \text{ even}}} (2^{\frac{n-l}{2}} \cdot g^\lambda)^2 + \sum_{\substack{\lambda^* \models n \\ n-l \text{ odd}}} 2 \cdot (2^{\frac{n-l-1}{2}} \cdot g^\lambda)^2. \end{aligned}$$

Conveniently, the powers of two in both summations turn out to be the same, and so

Corollary 3.2.2

$$n! = \sum_{\lambda^* \models n} 2^{n-l}(g^\lambda)^2.$$

To bring oscillating tableaux into the act, we need to talk about representations of the symplectic group (rather than the symmetric group). This discussion will be postponed until Section 7.

4 The Robinson-Schensted correspondence

Corollaries 3.1.7 and 3.2.2 were obtained from general theorems about group representations. However, the equations themselves can be viewed as purely combinatorial statements about tableaux. Hence it would be nice to have purely combinatorial (i.e., bijective) proofs of these results. The celebrated Robinson-Schensted correspondence [Rob 38, Sch 61] does exactly that. Although Robinson was the first to discover this algorithm, Schensted's form of the correspondence (discovered independently) is easier to understand. For that reason we will follow the latter's presentation.

4.1 Left-justified tableaux

We restate Corollary 3.1.7 for ease of reference.

Theorem 4.1.1

$$n! = \sum_{\lambda \vdash n} (f^\lambda)^2.$$

Combinatorial Proof. It suffices to find a bijection

$$\pi \overset{\text{R-S}}{\longleftrightarrow} (P, Q) \tag{1}$$

between permutations $\pi \in S_n$ and pairs of SYT P, Q of the same shape $\lambda \vdash n$. We first exhibit a map which, given a permutation, produces a tableaux pair.

$\underline{\pi \stackrel{\text{R–S}}{\to} (P, Q).}$ Suppose π is given in two-line form as

$$\pi = \begin{matrix} 1 & 2 & \cdots & n \\ x_1 & x_2 & \cdots & x_n \end{matrix} .$$

We will construct a sequence of tableaux

$$(P_0, Q_0) = (\phi, \phi);\ (P_1, Q_1);\ (P_2, Q_2);\ \cdots;\ (P_n, Q_n) = (P, Q) \tag{2}$$

where $x_1, x_2, \cdots, x_n$ will be *inserted* into the P's and $1, 2, \cdots, n$ will be *placed* in the Q's so that P_k and Q_k will have the same shape for all k. The operations of insertion and placement can be described as follows.

Suppose P is a *partial* tableau of shape μ, i.e., a filling of μ with a subset of the integers from 1 to n so that rows and columns increase. Let x be an element not in P. To *row insert x in P*, we use the following sequence of steps.

1. If x is bigger than every element of the first row of P, then put x at the end of that row ($p_{1,\mu_1+1} \leftarrow x$) and stop.

2. Otherwise, find the left-most element of the first row of P such that $p_{1,j_1} > x$ and replace this element by x (after storing its value for future use). We say that x *bumps* p_{1,j_1} from the first row.

3. Now iterate the first two steps. If p_{1,j_1} is bigger than every element in row 2 then put it at the end of the row and stop. Otherwise p_{1,j_1} replaces the left-most p_{2,j_2} larger than itself and this element is inserted into the third row, etc.

4. Since the p_{i,j_i} form an increasing sequence, at some point the algorithm must terminate with an element coming to rest at the end of some row.

As an example, suppose $x = 3$ and

$$P = \begin{matrix} 1\,2\,5\,8 \\ 4\,7 \\ 6 \end{matrix} .$$

To follow the path of the insertion of x into P, we will put elements that are bumped during the insertion in boldface type.

$$\begin{matrix} 1\,2\,5\,8 \ \leftarrow\mathbf{3} \\ 4\,7 \\ 6 \end{matrix} , \quad \begin{matrix} 1\,2\,\mathbf{3}\,8 \\ 4\,7 \ \leftarrow\mathbf{5} \\ 6 \end{matrix} , \quad \begin{matrix} 1\,2\,3\,8 \\ 4\,\mathbf{5} \\ 6 \ \leftarrow\mathbf{7} \end{matrix} , \quad \begin{matrix} 1\,2\,3\,8 \\ 4\,5 \\ 6\,\mathbf{7} \end{matrix} .$$

If row insertion of x into P yields partial tableau P' then we write $R_x(P) = P'$. It is easy to verify that P' will still have increasing rows and columns.

Placement of an element in a tableaux is an easy construction. Suppose that Q is a partial tableau of shape μ and that (i, j) is an *outer corner of* μ, meaning that $(i, j) \notin \mu$ but $\mu \cup (i, j)$ is the shape of a partition. If k is an integer, then the *placement of* k *in* Q *at cell* (i, j) is the tableau obtained by merely putting k in cell (i, j), i.e., $q_{i,j} := k$.

If we let

$$Q = \begin{matrix} 1 & 2 & 5 \\ 4 & 7 & \\ 6 & & \\ 8 & & \end{matrix} ,$$

then placing $k = 9$ in cell $(i, j) = (2, 3)$ yields

$$\begin{matrix} 1 & 2 & 5 \\ 4 & 7 & 9 \\ 6 & & \\ 8 & & \end{matrix} .$$

Clearly, if k is bigger than every element of Q then the array will remain a partial tableau.

We can finally describe how to build the sequence (2) from the permutation

$$\pi = \begin{matrix} 1 & 2 & \cdots & n \\ x_1 & x_2 & \cdots & x_n \end{matrix} .$$

Start with a pair of empty tableaux (P_0, Q_0). Assuming that (P_{k-1}, Q_{k-1}) have been constructed, define (P_k, Q_k) by

$$\begin{aligned} P_k &= R_{x_k}(P_{k-1}), \text{ and} \\ Q_k &= \text{the placement of } k \text{ into } Q_{k-1} \text{ at the cell } (i, j) \\ & \text{ where the row insertion terminates.} \end{aligned}$$

Note that the definition of Q_k insures that the shapes of P_k and Q_k are equal for all k. We call $P = P_n$ and $Q = Q_n$ the *P-tableau* and *Q-tableau* of π respectively.

We now give an example of the complete algorithm. Boldface numbers will be used to distinguish the elements of the upper line of π and hence also for the elements of the Q_k. Let

$$\pi = \begin{matrix} \mathbf{1} & \mathbf{2} & \mathbf{3} & \mathbf{4} & \mathbf{5} & \mathbf{6} \\ 2 & 4 & 3 & 6 & 5 & 1 \end{matrix} , \tag{3}$$

then the tableaux constructed by the algorithm are:

$$P_k: \quad \phi\,,\quad 2\,,\quad 2\,4\,,\quad \begin{matrix}2\,3\\4\end{matrix}\,,\quad \begin{matrix}2\,3\,6\\4\end{matrix}\,,\quad \begin{matrix}2\,3\,5\\4\,6\end{matrix}\,,\quad \begin{matrix}1\,3\,5\\2\,6\\4\end{matrix} \quad = P$$

$$Q_k: \quad \phi\,,\quad \mathbf{1}\,,\quad \mathbf{1\,2}\,,\quad \begin{matrix}\mathbf{1\,2}\\\mathbf{3}\end{matrix}\,,\quad \begin{matrix}\mathbf{1\,2\,4}\\\mathbf{3}\end{matrix}\,,\quad \begin{matrix}\mathbf{1\,2\,4}\\\mathbf{3\,5}\end{matrix}\,,\quad \begin{matrix}\mathbf{1\,2\,4}\\\mathbf{3\,5}\\\mathbf{6}\end{matrix} \quad = \mathbf{Q}.$$

Thus

$$\begin{matrix}\mathbf{1} & \mathbf{2} & \mathbf{3} & \mathbf{4} & \mathbf{5} & \mathbf{6}\\ 2 & 4 & 3 & 6 & 5 & 1\end{matrix} \stackrel{\text{R–S}}{\rightarrow} \left(\begin{matrix}1\,3\,5\\2\,6\\4\end{matrix}\quad,\quad \begin{matrix}\mathbf{1\,2\,4}\\\mathbf{3\,5}\\\mathbf{6}\end{matrix}\right).$$

To show that this map is a bijection, we create its inverse.

$\underline{(P,Q) \stackrel{\text{R–S}}{\rightarrow} \pi.}$ We merely reverse the above procedure step by step. Define $(P_n, Q_n) = (P,Q)$. Assuming that the pair (P_k, Q_k) has been constructed, we obtain x_k (the k^{th} element of π) and (P_{k-1}, Q_{k-1}) as follows.

Find the cell (i,j) containing the k in Q_k. Since this is the largest element in Q_k, $p_{i,j}$ must have been the last element to be bumped in the construction of P_k. Furthermore, the element that bumped it must be the right-most entry in row $i-1$ such that $p_{i-1,j_{i-1}} < p_{i,j}$. So replace $p_{i-1,j_{i-1}}$ by $p_{i,j}$ and find the entry of row $i-2$ that diplaced $p_{i-1,j_{i-1}}$, etc. Working back up the rows in this manner, we will finally remove an element p_{1,j_1} from the first row. Thus $x_k = p_{1,j_1}$, P_{k-1} is P_k after the deletion process described above is complete, and Q_{k-1} is Q_k with the k errased. Continuing in this way, we will eventually recover all the elements of π in reverse order. □

The Robinson-Schensted algorithm has many beautiful and surprising properties. The literature on this subject is so vast that we can only present a sampling of results here. The interested reader can consult the extensive bibliography in [Vie 84] for other sources.

4.1.1 Column insertion

One can obviously define *column insertion* of x into P by reversing the roles of rows and columns in the definition of insertion: x displaces the highest element of the first column of P larger than x, this element is bumped into the second column, etc. If the result of column inserting x into P is P' we write $C_x(P) = P'$. It turns out that the row and column insertion operators commute (operators should be read right to left).

Proposition 4.1.2 ([Sch 61]) *For any partial tableau P and positive integers $x, y \notin P$*

$$C_y R_x(P) = R_x C_y(P).$$

Proof. This proposition follows from the definitions of the two operators by an easy case-by-case argument. □

In the next result, π^r stands for the *reversal* of π, i.e., if $\pi = x_1 x_2 \cdots x_n$ then $\pi^r = x_n x_{n-1} \cdots x_1$.

Corollary 4.1.3 *If $\pi \stackrel{\text{R–S}}{\to} (P, Q)$. then $\pi^r \stackrel{\text{R–S}}{\to} (P^t, \cdot)$ where t denotes transpose.*

Proof . By definition, the P-tableau of π^r is

$$\begin{array}{rcll} R_{x_1} \cdots R_{x_{n-1}} R_{x_n}(\phi) & = & R_{x_1} \cdots R_{x_{n-1}} C_{x_n}(\phi) & \text{(initial tableau is empty)} \\ & = & C_{x_n} R_{x_1} \cdots R_{x_{n-1}}(\phi) & \text{(commutivity)} \\ & \vdots & & \\ & = & C_{x_n} C_{x_{n-1}} \cdots C_{x_1}(\phi) & \text{(induction)} \\ & = & P^t & \text{(def. of column insertion).} \ \Box \end{array}$$

4.1.2 The jeu de taquin

It would be nice to describe the Q-tableau for π^r. In order to do so, we must introduce a powerful operation of Schützenberger [Scü 63].

Suppose λ and μ are shapes such that $\lambda \subseteq \mu$. Then they form the *skew partition* $\lambda/\mu \stackrel{\text{def}}{=} \{(i,j) \in \lambda \mid (i,j) \notin \mu\}$. For example if

$$\lambda = \begin{array}{llll} \square & \square & \square & \square \\ \square & \square & \square & \\ \square & \square & \square & \end{array}$$

and

$$\mu = \begin{array}{ll} \square & \square \\ \square & \end{array}$$

then

$$\lambda/\mu = \begin{array}{llll} \bullet & \bullet & \square & \square \\ \bullet & \square & \square & \\ \square & \square & \square & \end{array}$$

where the missing boxes have been replaced by dots (black holes). *Skew tableaux* of both the standard and partial varieties are defined in the obvious way, filling the skew shape with an appropriate subset of the integers so that the rows and columns increase.

Now let Q be a skew tableau of shape λ/μ and let (i,j) be a corner cell of μ (called an *inner corner* of the skew shape). An (i,j)-*slide* is accomplished by performing the following sequence of operations.

1. If neither $(i,j+1)$ nor $(i+1,j)$ are in λ/μ then (i,j) is eliminated from the shape of Q and the algorithm terminates.

2. Otherwise, let $q_{i_1,j_1} = \min\{q_{i,j+1}, q_{i+1,j}\}$. (If one of the arguments of the min doesn't exist, then define it's value to be the element of Q which does appear.) Slide q_{i_1,j_1} into cell (i,j), creating a hole in position (i_1,j_1).

3. Now repeat the first two steps with (i,j) and (i_1,j_1) replaced by (i_1,j_1) and (i_2,j_2) respectively, etc. After a finite number of iterations, the hole will slide to the boundary of Q and be eliminated, at which point we stop.

If applying an (i,j)-slide to Q yields Q' then we write $\Delta_{i,j}Q = Q'$

To illustrate this procedure, let

$$Q = \begin{array}{cccc} \bullet & \bullet & 3 & 7 \\ \bullet & 1 & 4 & \\ 2 & 5 & 6 & \end{array} .$$

As usual boldface numbers are used to indicate the moving objects as we apply a (1,2)-slide:

$$\begin{array}{cccc} \bullet & \bullet & 3 & 7 \\ \bullet & 1 & 4 & \\ 2 & 5 & 6 & \end{array}, \quad \begin{array}{cccc} \bullet & \mathbf{1} & 3 & 7 \\ \bullet & \bullet & 4 & \\ 2 & 5 & 6 & \end{array}, \quad \begin{array}{cccc} \bullet & 1 & 3 & 7 \\ \bullet & \mathbf{4} & \bullet & \\ 2 & 5 & 6 & \end{array}, \quad \begin{array}{cccc} \bullet & 1 & 3 & 7 \\ \bullet & 4 & \mathbf{6} & \\ 2 & 5 & & \end{array} = \Delta_{1,2}Q.$$

Given an SYT Q we will build another SYT $S(Q)$ in the following manner. First construct a sequence of partial tableaux $Q = Q_n, Q_{n-1}, \cdots, Q_0 = \phi$. To get Q_{k-1} from Q_k, we first erase the element in cell $(1,1)$ of Q_k to form a tableau of shape $\mu/(1)$ for some μ . We let Q_{k-1} be the tableau obtained by applying $\Delta_{1,1}$ to this skew tableau. Finally, we put a k in cell (i,j) of $S(Q)$ if that was the box eliminated from the boundary when passing from Q_k to Q_{k-1}.

If we apply this algorithm to the Q-tableaux of the Robinson-Schensted example above, we obtain

$$
\begin{array}{llllllllllllll}
Q_k: & \begin{array}{l} 1\,2\,4 \\ 3\,5 \\ 6 \end{array} & , & \begin{array}{l} 2\,4 \\ 3\,5 \\ 6 \end{array} & , & \begin{array}{l} 3\,4 \\ 5 \\ 6 \end{array} & , & \begin{array}{l} 4 \\ 5 \\ 6 \end{array} & , & \begin{array}{l} 5 \\ 6 \\ \end{array} & , & \begin{array}{l} 6 \\ \\ \end{array} & , & \begin{array}{l} \phi \\ \\ \end{array}
\end{array}
$$

$$
\begin{array}{llllllllllllll}
S(Q): & \begin{array}{l} \bullet\,\bullet\,\bullet \\ \bullet\,\bullet \\ \bullet \end{array} & , & \begin{array}{l} \bullet\,\bullet\,6 \\ \bullet\,\bullet \\ \bullet \end{array} & , & \begin{array}{l} \bullet\,\bullet\,6 \\ \bullet\,5 \\ \bullet \end{array} & , & \begin{array}{l} \bullet\,4\,6 \\ \bullet\,5 \\ \bullet \end{array} & , & \begin{array}{l} \bullet\,4\,6 \\ \bullet\,5 \\ 3 \end{array} & , & \begin{array}{l} \bullet\,4\,6 \\ 2\,5 \\ 3 \end{array} & , & \begin{array}{l} 1\,4\,6\,. \\ 2\,5 \\ 3 \end{array}
\end{array}
$$

Furthermore, the reader can verify that the Q-tableau of $\pi^r = 1\ 5\ 6\ 3\ 4\ 2$ is just the transpose of $S(Q)$. This is not an accident.

Theorem 4.1.4 ([Scü 63]) *If* $\pi \stackrel{\text{R–S}}{\rightarrow} (P, Q)$ *then* $\pi^r \stackrel{\text{R–S}}{\rightarrow} (P^t, S(Q)^t)$. $\square$

Slides can also be used to prove another surprising theorem of Schützenberger.

Theorem 4.1.5 ([Scü 63]) *If* $\pi \stackrel{\text{R–S}}{\rightarrow} (P, Q)$ *then* $\pi^{-1} \stackrel{\text{R–S}}{\rightarrow} (Q, P)$.

Proof sketch. This result can also be demonstrated using a geometric form of the Robinson-Schensted correspondence due to Viennot [Vie 76]. Imagine π represented as a permutation matrix in the plane, i.e., the k^{th} element of π is represented by a point with cartesian coordinates (k, x_k). Suppose that the plane is illuminated from the origin so that each point of π casts a shadow whose boundaries are half-lines parallel to the coordinate axes. By reading this diagram from left to right, one obtains a picture of the Robinson-Schensted algorithm as if on a time line (the k^{th} insertion takes place as we pass the line $x = k$). One can read off the entries of the P- and Q-tableaux as certain coordinates on the y- and x-axes respectively. Once this is established, the theorem is immediate since passing from π to π^{-1} merely interchanges the two axes. $\square$

Two more definitions are needed before we will be able to define the 'jeu de taquin' (or 'teasing game'). An *anti-diagonal strip* is the skew shape consisting of the cells $(n+1, 1)$; $(n, 2)$; $\cdots$; $(1, n+1)$. If $\pi = x_1 x_2 \cdots x_n$ then the corresponding *anti-diagonal strip tableau* has x_j in column j. For example, $\pi = 1\,4\,3\,2$ corresponds to the tableau

$$
\begin{array}{llll}
\bullet & \bullet & \bullet & 2 \\
\bullet & \bullet & 3 & \\
\bullet & 4 & & \\
1 & & &
\end{array}
\quad .
$$

Now given an anti-diagonal strip tableaux, we can play *jeu de taquin.* Start by choosing any inner corner (i, j) and applying $\Delta_{i,j}$. Now choose any inner corner

(i', j') of the new skew shape and apply $\Delta_{i',j'}$, etc., until we get a left-justified (non-skew) SYT. One possible game that could be played on the tableau above is

$$\begin{array}{llll} \bullet & \bullet & \bullet & 2 \\ \bullet & \bullet & 3 & \\ \bullet & 4 & & \\ 1 & & & \end{array} \overset{\Delta_{3,1}}{\longrightarrow} \begin{array}{llll} \bullet & \bullet & \bullet & 2 \\ \bullet & \bullet & 3 & \\ 1 & 4 & & \end{array} \overset{\Delta_{1,3}}{\longrightarrow} \begin{array}{lll} \bullet & \bullet & 2 \\ \bullet & \bullet & 3 \\ 1 & 4 & \end{array} \overset{\Delta_{2,2}}{\longrightarrow} \begin{array}{lll} \bullet & \bullet & 2 \\ \bullet & 3 & \\ 1 & 4 & \end{array} \overset{\Delta_{2,1}}{\longrightarrow} \begin{array}{lll} \bullet & \bullet & 2 \\ 1 & 3 & \\ 4 & & \end{array} \overset{\Delta_{1,2}}{\longrightarrow} \begin{array}{ll} \bullet & 2 \\ 1 & 3 \\ 4 & \end{array} \overset{\Delta_{1,1}}{\longrightarrow} \begin{array}{ll} 1 & 2 \\ 3 & \\ 4 & \end{array}.$$

If applying jeu de taquin to the anti-diagonal strip of π yields a SYT P we write $J(\pi) = P$. It is not clear that the operation J is well-defined. However,

Theorem 4.1.6 ([Scü 76]) *The tableau $J(\pi)$ is independent of the choice of inner corners made while playing the game. Furthermore, if P is the P-tableau of π then $J(\pi) = P$.*

Proof sketch. This proof is due to Thomas [Tho 77] .Showing that $J(\pi)$ does not depend on the order in which the inner corners are filled is a delicate case-by-case argument. Once this is established, a beautiful connection between the jeu de taquin and the Robinson-Schensted map appears. Let's choose to fill all the corners in a given row from right to left, starting with the lowest row and working up. After $k-1$ rows have been filled, the portion of the array in these rows will be the partial tableau P_k of the sequence (2). From this, it is easy to see that filling the k^{th} row from the bottom is equivalent to the row insertion of x_{k+1} into P_k. Hence the lower portion of the array must now contain P_{k+1} and induction yields the fact that $J(\pi) = P$. $\square$

4.1.3 Increasing and decreasing subsequences

One of Schensted's original motivations for constructing his map was to investigate lengths of increasing and decreasing subsequences of a permutation $\pi = x_1 x_2 \cdots x_n$. An *increasing subsequence of* π is a subsequence $x_{i_1} x_{i_2} \cdots x_{i_k}$ such that $x_{i_1} < x_{i_2} < \cdots < x_{i_k}$. *Decreasing subsequences* are similarly defined. For example, 2 3 6 and 4 3 1 are respectively increasing and decreasing subsequences of the permutaion listed as equation (3).

Theorem 4.1.7 ([Sch 61]) *The length of a longest increasing subsequence of π is the length of the first row of it's P-tableau. The length of a longest decreasing subsequence of π is the length of the first column.*

Proof sketch. For increasing subsequences, one can inductively prove a stronger result, viz., if k enters P_{k-1} in column j then the length of the longest increasing subsequence ending in x_k must be j. The statement for decreasing subsequences now follows from Corollary 4.1.3. □

The reader should note that while the length of P's first row is the length of a longest increasing subsequence of π, the elements themselves do not form an increasing subsequence. This can be verified using our running example. It is possible, with a little more care, to recover an increasing subsequence of maximum length from the Robinson-Schensted algorithm.

Greene [Gre 74] has generalized this theorem to other types of subsequences. A subsequence of π is called *k-increasing* (respectively *k-decreasing*) if it is the union of k increasing (respectively decreasing) subsequences. Thus a 1-increasing subsequence is merely an increasing one. The subsequence $2\ 4\ 3\ 6\ 5\ = 2\ 3\ 6 \cup 4\ 5$ of the permutation in (3) is 2-increasing but not 1-increasing, while π itself is 3-increasing.

Theorem 4.1.8 ([Gre 74]) *The length of a longest k-increasing (respectively k-decreasing) subsequence of π is the sum of the lengths of the first k rows (respectively columns) of it's P-tableau.*

Proof sketch. Given a tableau P, we define the *row word* of P to be the permutation obtained by reading of the rows of P from left to right, starting with the last row and moving up. Our running example has row word $4\ 2\ 6\ 1\ 3\ 5$. It is easy to see that if P has row word π then

1. the P-tableau of π is P itself, and
2. the first k rows of P form a k-increasing subsequence of π of maximum length.

Hence the theorem is true for permutations that are row words. To show that it holds in general, Greene proves that any permutation with P-tableau P can be transformed into P's row word by a sequence of adjacent transpositions (the so-called *Knuth transpositions* [Knu 70]) which leave both the P-tableau and the maximum length of a k-increasing subsequence invariant. □

4.2 Shifted tableaux

Suppose λ^* is a strict partition of n having length l. The *main diagonal* of the shape of λ^* consists of the cells (i, i) where $1 \leq i \leq l$. All other cells are *off-diagonal*, so the number of off-diagonal cells is $n - l$. This fact will be useful in the combinatorial proof of Corollary 3.2.2 which we restate here.

Theorem 4.2.1

$$n! = \sum_{\lambda^* \models n} 2^{n-l} (g^\lambda)^2.$$

Combinatorial proof. It suffices to find a bijection

$$\pi \longleftrightarrow (P^*, Q^*) \tag{4}$$

between permutations $\pi \in S_n$ and pairs of SST P^*, Q^* of the same shape $\lambda^* \models n$ where Q^* has a subset of it's off-diagonal elements distinguished in some manner. We will distinguish an element k by writing k'. Sagan [Sag 79] was the first to find such a map, but his correspondence did not have many of the properties of the original Robinson-Schensted algorithm. Later Sagan [Sag 87] and Worley [Wor 84] independently found a better bijection that does enjoy most of these properties. It is this version that we present.

To construct the map from permutations to tableau pairs, we create a sequence of shifted partial tableau pairs analogous to the sequence (2). Thus we need only discuss the analogs of insertion and placement for shifted tableaux.

To insert x_k into a partial tableau P^*_{k-1}, we start row inserting x_k as usual. If a diagonal element is never displaced, then insertion stops with an element coming to rest at the end of some row. This is called a *Schensted* insertion. If, on the other hand, some $p^*_{i,i}$ is bumped, then insert it in *column* $i+1$. Continue column inserting until an element comes to rest at the end of some column. This type of insertion is called *non-Schensted*. In either case, let (i,j) be the cell filled by the last bump.

Q^*_k is obtained by placing an element in cell (i,j) of Q^*_{k-1}. This element will be either a k if the insertion was Schensted, or a k' if the insertion was non-Schensted. Since a non-Schensted insertion can never terminate on the main diagonal, all primed elements will be off-diagonal.

It is an easy matter to construct the inverse map, since the distinguished elements in Q^* indicate whether to start the deletion process by rows or columns. The details are left to the reader. □

If we compute the shifted tableaux associated with the permutation $\pi = 2\ 1\ 6\ 5\ 4\ 3$, we get the sequence

$$\begin{array}{lllllllll} P^*_k: & \phi\,, & 2\,, & 1\,2\,, & 1\,2\,6\,, & \begin{array}{l}1\,2\,5\\ \;\;\;\;6\end{array}, & \begin{array}{l}1\,2\,4\\ \;\;\;5\,6\end{array}, & \begin{array}{l}1\,2\,3\,6\\ \;\;\;4\,5\end{array} & = P^* \end{array}$$

$$\begin{array}{lllllllll} Q^*_k: & \phi\,, & 1\,, & 1\,2'\,, & 1\,2'\,3\,, & \begin{array}{l}1\,2'\,3\\ \;\;\;\;4\end{array}, & \begin{array}{l}1\,2'\,3\\ \;\;\;4\,5'\end{array}, & \begin{array}{l}1\,2'\,3\,6'\\ \;\;\;4\,5'\end{array} & = Q^*. \end{array}$$

We should mention that MacLarnan [MaL 86] has found a way to construct other bijections between permutations and tableaux pairs using recursions satisfied by f^λ g^λ, and $\tilde{f}^\lambda_k$. All of these maps have the property that inverting the permutation interchanges the tableaux, which is *not* true for the algorithm above. However, Haiman [Hai pr] has developed a procedure called *mixed insertion* which does interchange the outputs of the Sagan-Worely algorithm when applied to π^{-1}.

5 Tableaux with repetitions

Many of the results of the first four sections can be generalized to tableaux with repeated entries. First we must define such arrays precisely.

5.1 Generalized Young tableaux

A *generalized Young tableau* (GYT), T, of shape λ is a filling of the shape with positive integers such that the rows weakly increase and the columns strictly increase. These arrays are also called *semi-standard tableaux* or *column strict reverse plane partitions* (the term 'reverse' is an historical accident coming from the fact that partitions were usually listed with parts in *decreasing* order). One possible GYT of shape $(4, 4, 1)$ is

$$\begin{array}{llll} 1 & 1 & 2 & 2 \\ 2 & 3 & 4 & 4 \\ 4 & & & \end{array}$$

Note that if $\mu = (\mu_1, \mu_2, \cdots, \mu_l)$ is a partition of n then we can also write $\mu = (1^{m_1}, 2^{m_2}, \cdots, n^{m_n})$ where m_k is the number of parts of μ equal to k. The same notation applies if the parts of μ are arranged in a GYT, T. In this case, μ is called the *content of* T. For example, the tableau above has content $(1^2, 2^3, 3^1, 4^3)$. Using the set of variables $\mathbf{x} = \{x_1, x_2, \cdots, x_n\}$ we can associate with T a *monomial* $m(T) = x_1^{m_1} x_2^{m_2} \cdots x_n^{m_n}$. The monomial of the tableau above is $m(T) = x_1^2 x_2^3 x_3^1 x_4^3$.

Finally, if λ is a partition of length at most n, we define the corresponding *Schur function* to be

$$s_\lambda(\mathbf{x}) = \sum_{T \in \mathcal{T}_\lambda(n)} m(T)$$

where $\mathcal{T}_\lambda(n)$ is the set of all GYT of shape λ and entries of size at most n. If we let $\lambda = (2, 1)$ and $n = 3$, then the tableaux in $\mathcal{T}_{(2,1)}(3)$ are

$$\begin{array}{l} 1\,1 \\ 2 \end{array}, \quad \begin{array}{l} 1\,2 \\ 2 \end{array}, \quad \begin{array}{l} 1\,1 \\ 3 \end{array}, \quad \begin{array}{l} 1\,3 \\ 3 \end{array}, \quad \begin{array}{l} 2\,2 \\ 3 \end{array}, \quad \begin{array}{l} 2\,3 \\ 3 \end{array}, \quad \begin{array}{l} 1\,2 \\ 3 \end{array}, \quad \begin{array}{l} 1\,3 \\ 2 \end{array}$$

and so

$$s_{(2,1)}(x_1, x_2, x_3) = {x_1}^2 x_2 + x_1 {x_2}^2 + {x_1}^2 x_3 + x_1 {x_3}^2 + {x_2}^2 x_3 + x_2 {x_3}^2 + 2x_1 x_2 x_3 \ .$$

Clearly if λ is a partition of n, then the coefficient of the monomial $x_1 x_2 \cdots x_n$ in $s_\lambda(x_1, x_2, \cdots, x_n)$ is just f^λ. It is also true (although it is not obvious from our definition) that $s_\lambda(\mathbf{x})$ is a *symmetric function*, i.e., permuting the variables in $\mathbf{x}$ does not change the polynomial s_λ. We will have more to say about symmetric functions shortly.

5.2 Generalized shifted tableaux

Let λ^* be a strict partition. A *generalized shifted tableau* (GST), T^* , of shape λ^* is obtained by filling the shifted shape with elements from the totally ordered alphabet $\mathcal{A}' = \{1' < 1 < 2' < 2 < 3' < 3 < \cdots\}$ so that

1. T^* is weakly increasing along rows and columns and strictly increasing along diagonals, and
2. for every integer k, there is at most one k' in each row and at most one k in each column of T^* .

An example of such a tableau is

$$T^* = \begin{array}{cccccccccc} 1' & 1 & 1 & 2' & 3' & 3 & 3 & 4' \\ & 2 & 2 & 2 & 3' & 4' & 4 & 4 \\ & & 3' & 3 & 3 & & & \\ & & & 4 & 4 & & & \end{array} .$$

Note that conditions 1 and 2 imply that, for fixed k, the cells occupied by all the elements of the form k or k' form a union of skew hooks. A *skew hook* is a skew shape that is connected (i.e., one can travel from one cell to any other by passing through cells adjacent by an edge) and contains at most one cell on every diagonal. The elements 4 and $4'$ in T^* above lie in the union of two skew hooks while all other fixed integers and their primes lie in only one. Furthermore, condition 2 determines the character (primed or not) of each element in a skew hook, except at the lower left-hand end. This observation will be important when we discuss the analog of Knuth's algorithm for generalized shifted tableaux.

If T^* is a GST with $t^*_{i,j} \leq n$ for all $(i,j) \in \lambda^*$ then it's associated monomial is $m(T^*) = x_1^{m_1} x_2^{m_2} \cdots x_n^{m_n}$ where m_k is the number of entries of T^* equal to k or k'.

Our example tableau has monomial $m(T^*) = {x_1}^3{x_2}^4{x_3}^7{x_4}^6$. Associated with each strict partition λ^* , $l(\lambda^*) \le n$, is a *Schur Q-function* defined by

$$Q_\lambda(\mathbf{x}) = \sum_{T^* \in \mathcal{T}^*_\lambda(n)} m(T^*),$$

$\mathcal{T}^*_\lambda(n)$ being the set of all GST with shape λ^* and entries of size at most n. Like $s_\lambda(\mathbf{x})$, $Q_\lambda(\mathbf{x})$ is a symmetric function. However, because of the presence of primed elements, the coefficient of $x_1 x_2 \cdots x_n$ in $Q_\lambda(\mathbf{x})$ for $\lambda^* \models n$ turns out to be $2^n g^\lambda$.

5.3 Symplectic tableaux

Consider the alphabet $\bar{\mathcal{A}} = \{1 < \bar{1} < 2 < \bar{2} < 3 < \bar{3} < \cdots\}$. A *symplectic tableau* (SPT), $\tilde{T}$, of shape $\lambda \vdash n$ is a GYT with entries from $\bar{\mathcal{A}}$ satisfying the extra constraint

$$\text{for all } i \le l(\lambda), \text{ the elements in row } i \text{ are all of size at least } i. \tag{5}$$

Equation (5) is called the *symplectic condition.* An example of such an array is

$$\tilde{T} = \begin{array}{cccccc} 1 & \bar{1} & 2 & 2 & 2 & \bar{3} \\ 2 & 2 & \bar{2} & \bar{2} & \bar{2} & \\ \bar{3} & \bar{3} & & & & \end{array} .$$

For SPT we use the set of variables $\mathbf{x}^{\pm 1} = \{x_1, {x_1}^{-1}, x_2, {x_2}^{-1}, \cdots, x_n, {x_n}^{-1}\}$. The monomial of $\tilde{T}$ is given by

$$m(\tilde{T}) = {x_1}^{m_1}{x_1}^{-\bar{m}_1}{x_2}^{m_2}{x_2}^{-\bar{m}_2} \cdots x_n^{m_n}{x_n}^{-\bar{m}_n}$$

where m_k (respectively $\bar{m}_k$) is the number of k's (respectively $\bar{k}$'s) in $\tilde{T}$. The SPT above has $m(\tilde{T}) = {x_1}^1{x_1}^{-1}{x_2}^5{x_2}^{-3}{x_3}^{-3}$. The *symplectic Schur function* associated with a partition λ of length at most n is

$$sp_\lambda(\mathbf{x}^{\pm 1}) = \sum_{\tilde{T} \in \tilde{\mathcal{T}}_\lambda(n)} m(\tilde{T})$$

where the reader will already have guessed that $\tilde{\mathcal{T}}_\lambda(n)$ is the set of all SPT of shape λ with entries of size at most n. The symplectic Schur function is symmetric in the variables $\mathbf{x}^{\pm 1}$.

6 Enumeration of generalized tableaux

We will now present analogs of the hook and determinantal formulae for generalized tableaux.

6.1 The ordinary case

Let T be a GYT of shape λ , then we say T *partitions* m if $\sum_{(i,j)\in\lambda} t_{i,j} = m$. Letting $p_\lambda(m)$ be the number of such tableaux, we have the following generating function analog of the hook formula.

Theorem 6.1.1

$$\sum_{m\geq 0} p_\lambda(m)x^m = x^{N(\lambda)} \prod_{(i,j)\in\lambda} \frac{1}{1-x^{h_{i,j}}}$$

where $N(\lambda) = \sum_{i\geq 1} i\lambda_i$.

Proof. Stanley [Sta 72] was the first to prove this using his theory of poset partitions. We will present a beautiful bijective proof of Hillman and Grassl [H-G 76].

A *reverse plane partition* is like a GYT except that the columns need only weakly increase and 0 is allowed as an array entry. Let $r_\lambda(m)$ be the number of reverse plane partitions of m. There is a simple bijection between GYT and reverse plane partitions of the same shape λ . Merely take the GYT and subtract 1 from every element of the first row, 2 from every element of the second, etc. Since this takes away a total of $N(\lambda)$ from each GYT, it suffices to prove that

$$\sum_{m\geq 0} r_\lambda(m)x^m = \prod_{(i,j)\in\lambda} \frac{1}{1-x^{h_{i,j}}}. \tag{6}$$

The right hand side of this equation counts (linear) partitions ν all of whose parts come from the multiset (i.e., a set with repeated elements) $\{h_{i,j} | (i,j)\in\lambda\}$. Thus we need a bijection

$$T \longleftrightarrow \nu$$

between GYT's T and partitions $\nu = (\nu_1, \cdots, \nu_l)$ whose parts are all of the form $h_{i,j}$ such that $\sum_{(i,j)\in\lambda} t_{i,j} = \sum_{k\geq 1} \nu_k$.

$\underline{T \to \nu\ .}$ Given T, we will produce a sequence of reverse plane partitions

$$T = T_0,\ T_1,\ T_2,\ \cdots, T_f = \text{ tableau of 0's}$$

where T_k will be obtained from T_{k-1} by subtracting one from all the elements of a skew hook H of T_k such that $|H| = h_{a,c}$ for some $(a,c)\in\lambda$. The cells of the skew hook are defined recursively as follows.

Let (a,b) be the right-most highest cell of T containing a non-zero element. Then

$$(a,b)\in H \text{ and if } (i,j)\in H \text{ then} \begin{cases} (i,j-1)\in H & \text{if } t_{i,j-1} = t_{i,j} \\ (i+1,j)\in H & \text{otherwise} \end{cases}$$

i.e., move down unless forced to move left so as not to violate the weakly increasing condition along the rows (once the ones are subtracted). Continue this process until the induction rule fails. At this point we must have stopped at the end of some column, say column c. It is easy to see that after subtracting one from the elements in H, the array remains a reverse plane partition and the amount subtracted is $h_{a,c}$.

As an example, let

$$T = \begin{array}{cccc} 1 & 2 & 2 & 2 \\ 3 & 3 & 3 & \\ 3 & & & \end{array} .$$

Then the skew hook H consists of the diamonds in the shape

$$\begin{array}{cccc} \square & \diamond & \diamond & \diamond \\ \diamond & \diamond & \square & \\ \diamond & & & \end{array} .$$

After subtraction, we have

$$T_1 = \begin{array}{cccc} 1 & 1 & 1 & 1 \\ 2 & 2 & 3 & \\ 2 & & & \end{array} .$$

To obtain the rest of the T_k, we iterate this process. The complete list for our example array, together with the corresponding $h_{i,j}$, is

$$\begin{array}{llllll l} T_k : & \begin{array}{l} 1\,2\,2\,2 \\ 3\,3\,3 \\ 3 \end{array}, & \begin{array}{l} 1\,1\,1\,1 \\ 2\,2\,3 \\ 2 \end{array}, & \begin{array}{l} 0\,0\,0\,0 \\ 1\,2\,3 \\ 1 \end{array}, & \begin{array}{l} 0\,0\,0\,0 \\ 1\,2\,2 \\ 1 \end{array}, & \begin{array}{l} 0\,0\,0\,0 \\ 1\,1\,1 \\ 1 \end{array}, & \begin{array}{l} 0\,0\,0\,0 \\ 0\,0\,0 \\ 0 \end{array}. \\ & & & & & & \\ h_{i,j} : & h_{1,1} & h_{1,1} & h_{2,3} & h_{2,2} & h_{2,1} . & \end{array}$$

Hence $\nu = (h_{1,1}, h_{1,1}, h_{2,3}, h_{2,2}, h_{2,1})$.

$\underline{\nu \to T.}$ Now given a partition of hooklengths ν we must rebuild T. First, however, we must know in what order the hooklengths were removed. It is easy to see that if $h_{i,j}, h_{i',j'} \in \nu$ then $h_{i,j}$ was removed before $h_{i',j'}$ if an only if $i < i'$, or $i = i'$ and $j \geq j'$. Once this is established reversing the subtraction process is straight-forward. For details the reader can consult [H-G 76] . $\square$

We can use Theorem 6.1.1 to derive the hook formula. It follows from general facts about poset partitions that if λ is a partition of n then

$$\sum_{m \geq 0} r_\lambda(m) x^m = \frac{p(x)}{\prod_{k=1}^{n}(1 - x^k)}$$

where $p(x)$ is a polynomial in x such that $p(1) = f^\lambda$. Combining this with equation (6) we obtain

$$p(x) = \frac{\prod_{k=1}^{n}(1-x^k)}{\prod_{(i,j)\in\lambda}(1-x^{h_{i,j}})}.$$

Now taking the limit as x approaches 1 yields $f^\lambda = n!/\prod_{(i,j)\in\lambda} h_{i,j}$.

We saw in Section 4.1.2 that the Robinson-Schensted map and the jeu de taquin are equivalent (Theorems 4.1.5 and 4.1.6). Kadell [Kad pr] has shown that the Hillman-Grassl algorithm is just another form of the jeu de taquin. Hence all three constructs are really the same. (Gansner [Gan 78] also noted this for the special case of rectangular arrays in his thesis.)

Before discussing the analog of the determinantal formula, we must talk briefly about the theory of symmetric functions. A *symmetric function* in the variables $\mathbf{x} = (x_1, x_2, \cdots, x_n)$ is a polynomial $f(\mathbf{x})$ with coefficients in $\mathbf{C}$ which is invariant under permutation of variables, i.e., for all $\pi \in S_n$ we must have $f(x_{\pi(1)}, x_{\pi(2)}, \cdots, x_{\pi(n)}) = f(x_1, x_2, \cdots, x_n)$. The set of all symmetric functions in n variables forms an algebra denoted Λ_n.

There are several well-known bases for Λ_n. The obvious one consists of the polynomials obtained by symmetrizing a given monomial. Specifically the *monomial symmetric function* corresponding to a partition $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_n)$ (where we permit λ to have parts equal to 0) is the polynomial

$$m_\lambda(\mathbf{x}) = \sum_{\pi\in S_n} x_{\pi(1)}^{\lambda_1} x_{\pi(2)}^{\lambda_2} \cdots x_{\pi(n)}^{\lambda_n}.$$

For example

$$m_{(2,1)}(x_1, x_2, x_3) = {x_1}^2 x_2 + x_1 {x_2}^2 + {x_1}^2 x_3 + x_1 {x_3}^2 + {x_2}^2 x_3 + x_2 {x_3}^2.$$

The three other important families of symmetric functions are as follows.

1. The k^{th} *elementary symmetric function* defined by

$$e_k(\mathbf{x}) = m_{(1^k)}(\mathbf{x}) = \sum_{i_1<\cdots<i_k} x_{i_1} \cdots x_{i_k}.$$

2. The k^{th} *power sum symmetric function* defined by

$$p_k(\mathbf{x}) = m_{(k)}(\mathbf{x}) = \sum_{i\geq 1} x_i^k.$$

3. The k^{th} *complete homogeneous symmetric function* defined by

$$h_k(\mathbf{x}) = \sum_{\lambda\vdash k} m_\lambda(\mathbf{x}) = \sum_{i_1\leq\cdots\leq i_k} x_{i_1} \cdots x_{i_k}.$$

By way of illustration, when $k = 2$ and $n = 3$:

$$\begin{aligned} e_2(x_1, x_2, x_3) &= x_1x_2 + x_1x_3 + x_2x_3, \\ p_2(x_1, x_2, x_3) &= {x_1}^2 + {x_2}^2 + {x_3}^2, \\ h_2(x_1, x_2, x_3) &= {x_1}^2 + {x_2}^2 + {x_3}^2 + x_1x_2 + x_1x_3 + x_2x_3. \end{aligned}$$

We can extend these definitions to partitions λ by letting

$$\begin{aligned} e_\lambda(\mathbf{x}) &= e_{\lambda_1}(\mathbf{x})e_{\lambda_2}(\mathbf{x}) \cdots e_{\lambda_n}(\mathbf{x}), \\ p_\lambda(\mathbf{x}) &= p_{\lambda_1}(\mathbf{x})p_{\lambda_2}(\mathbf{x}) \cdots p_{\lambda_n}(\mathbf{x}), \text{ and} \\ h_\lambda(\mathbf{x}) &= h_{\lambda_1}(\mathbf{x})h_{\lambda_2}(\mathbf{x}) \cdots h_{\lambda_n}(\mathbf{x}). \end{aligned}$$

In the next result, the *length* of a partition will be the number of *non-zero* parts.

Theorem 6.1.2 *The following sets are all bases for* Λ_n

1. $\{m_\lambda \mid l(\lambda) \le n\}$,
2. $\{e_\lambda \mid l(\lambda) \le n\}$,
3. $\{p_\lambda \mid l(\lambda) \le n\}$, *and*
4. $\{h_\lambda \mid l(\lambda) \le n\}$. $\square$

The proof of Theorem 6.1.2 can be found in any book on symmetric functions, e.g., [Mad 79]

The Jacobi-Trudi identity is the Schur function analog of the determinantal formula.

Theorem 6.1.3 *If* $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_n)$ *then*

$$s_\lambda(\mathbf{x}) = \det[h_{\lambda_i - i + j}(\mathbf{x})]$$

where the determinant is $n \times n$. $\square$

Both the proofs that we gave of the determinantal formula can be generalized to prove this Theorem. In particular, weighting the lattice paths of Gessel and Viennot appropriately results in a combinatorial proof.

The Jacobi-Trudi formula also has a dual version using elementary symmetric functions.

Theorem 6.1.4 *If* $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_n)$ *then*

$$s_{\lambda^t}(\mathbf{x}) = \det[e_{\lambda_i - i + j}(\mathbf{x})]_{n \times n}$$

where λ^t *is the transpose of the shape* λ *(also called the conjugate)* $\square$

6.2 The shifted case

A *shifted reverse plane partition* is defined in exactly the same way as an ordinary one, only using a shifted shape. Let $q_\lambda(m)$ be the number of shifted reverse plane partitions of m having shape λ^*. Then we have the following analog of Theorem 6.1.1

Theorem 6.2.1

$$\sum_{m\geq 0} q_\lambda(m)x^m = \prod_{(i,j)\in\lambda^*} \frac{1}{1-x^{h^*_{i,j}}}. \ \Box$$

Gansner [Gan 78] was the first to prove Theorem 6.2.1. He used generating function manipulations to obtain the shifted result from facts about symmetric (left-justified) reverse plane partitions. Later, Sagan [Sag 82] gave a bijective proof based on the Hillman-Grassl algorithm. He also showed that similar techniques yield many other product generating function identities.

6.3 The symplectic case

A symplectic analog of the Jacobi-Trudi identity can be derived from the Weyl character formula, a deep result in the representation theory of Lie groups. In what follows, $\tilde{h}_k(\mathbf{x}^{\pm 1}) \stackrel{\text{def}}{=} h_k(x_1, {x_1}^{-1}, x_2, {x_2}^{-1}, \cdots, x_n, {x_n}^{-1})$.

Theorem 6.3.1 *If* $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_n)$ *then*

$$sp_\lambda(\mathbf{x}^{\pm 1}) = \frac{1}{2}\det[\tilde{h}_{\lambda_i - i + j}(\mathbf{x}^{\pm 1}) + \tilde{h}_{\lambda_i - i - j + 2}(\mathbf{x}^{\pm 1})]. \ \Box$$

7 Characters of representations

7.1 Ordinary characters

Let V be an n-dimensional G-module. By picking a basis for V we can view the corresponding representation ρ as a homomorphism from G to GL_n, the group of all $n \times n$ matrices over $\mathbf{C}$. This viewpoint will be useful in our discussion of group characters.

If $\rho : G \to GL_n$ is a representation then its *character* is the map $\chi : G \to \mathbf{C}$ defined by $\chi(g) = tr\rho(g)$ for all $g \in G$. If V is a G-module for ρ we say that V *affords* χ. Since the trace function is invariant under change of basis, $\chi(g)$ is well-defined. Furthermore, if g_1 and g_2 are conjugate in G, then $g_1 = hg_2h^{-1}$ for some $h \in G$ and so $\rho(g_1) = \rho(h)\rho(g_2)\rho(h)^{-1}$. Thus $\chi(g_1) = \chi(g_2)$ since similar matrices have the same trace. This means that χ is a *class function*, i.e., a function constant on conjugacy classes of G.

Let us look at a some examples. If $\rho : S_n \to GL_n$ is the defining representation, then $\rho(\pi)$ is the permutation matrix of $\pi \in S_n$. Thus $\chi(\pi)$ is just the number of fixed-points of π. Now let $G = \{g_1, g_2, \cdots, g_n\}$ be any group and let ρ be the regular representation which has the group algebra $\mathbf{C}(G)$ as module. Thus for any $g \in G$, $\chi(g)$ is the number of fixed-points of g acting on the basis $\vec{g}_1, \vec{g}_2, \cdots, \vec{g}_n$. It follows that

$$\chi(g) = \begin{cases} n & \text{if } g = e \\ 0 & \text{if } g \neq e \end{cases}$$

where e is the identity element of G. Finally note that for any G-module V, $\rho(e)$ is the identity matrix so $\chi(e) = \dim V$.

Recall that the irreducible representations of S_n are given by the Specht modules S^λ for $\lambda \vdash n$. If χ^λ is the corresponding character, then we know that $\chi^\lambda(e) = f^\lambda$. To describe the rest of the character values, we will use the notation χ^λ_μ for the value of χ^λ on the class of permutations of cycle-type $\mu \vdash n$.

Theorem 7.1.1 *If $\lambda \vdash n$ then*

$$s_\lambda(\mathbf{x}) = \frac{1}{n!} \sum_{\mu \vdash n} c_\mu \; \chi^\lambda_\mu \; p_\lambda(\mathbf{x})$$

where c_μ is the number of elements of S_n in the class μ. $\square$

Hence the Schur function $s_\lambda(\mathbf{x})$ is just the cycle index generating function (in the sense of Pólya theory) for the character of the corresponding Specht module.

Now let us consider representations of matrix groups. If G is a group of matrices, then $\rho : G \to GL_n$ is called a *polynomial representation* if, for every $X \in G$, the entries of the matrix $\rho(X)$ are polynomials in the entries of X. As examples, the trivial representation is clearly polynomial. The identity map $id : GL_n \to GL_n$ is a polynomial representation, called the *defining representation.* Also the determinant $\det : GL_n \to GL_1$ is a representation which is polynomial. It follows from the work of Schur [Scu 01] that polynomial representations are very nice.

Theorem 7.1.2 *Polynomial representations of GL_n are completely reducible.* $\square$

Lest the reader get the idea that every representation is completely reducible, consider the representation of GL_n defined by

$$\rho(X) = \begin{bmatrix} 1 & \log|\det X| \\ 0 & 1 \end{bmatrix}$$

for all $X \in GL_n$. The x-axis is an invariant subspace, so if ρ were completely reducible it would have two decompose as the direct sum of two invariant one-dimensional subspaces. But this would mean that there would exist a fixed matrix Y such that

$$Y\rho(X)Y^{-1} = \left[\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array}\right]$$

for every $X \in GL_n$ (since both eigenvalues of $\rho(X)$ equal 1) which is absurd.

Now let $\rho : GL_n \to GL(V)$ be a polynomial representation with character χ. Let $X \in GL_n$ be diagonalizable with eigenvalues $\mathbf{x} = \{x_1, x_2, \cdots, x_n\}$ and let corresponding diagonal matrix be $\mathrm{diag}(x_1, x_2, \cdots, x_n)$. Since χ is a class function, $\chi(X) = \chi(\mathrm{diag}(x_1, x_2, \cdots, x_n))$ which is a polynomial in $x_1, x_2, \cdots, x_n$. Since the diagonalizable matrices are dense in GL_n and ρ is continuous (being polynomial) it follows that $\chi(X)$ is a polynomial in the eigenvalues of X for any $X \in GL_n$. Furthermore this polynomial must be a symmetric function of $x_1, x_2, \cdots, x_n$ (since permuting the elements in a diagonal matrix leaves one in the same conjugacy class). As an example, note that if ρ is the defining representation then its character is $\chi(X) = x_1 + x_2 + \cdots + x_n$.

It is natural to ask which symmetric functions give the characters of the irreducible polynomial GL_n-modules. Again, the Schur functions play a role.

Theorem 7.1.3 ([Scu 01]) *The irreducible polynomial representations of GL_n are indexed by partitions λ of length at most n. If λ is such a partition with corresponding module V^λ then the character afforded by V^λ is*

$$\phi^\lambda(X) = s_\lambda(\mathbf{x})$$

where $\mathbf{x} = \{x_1, x_2, \cdots, x_n\}$ is the set of eigenvalues of X. $\square$

If V is any GL_n-module, then so is the k^{th} tensor power $V^{\otimes k}$ since we have the natural action

$$X(\vec{v}_1 \otimes \vec{v}_2 \otimes \cdots \otimes \vec{v}_k) \stackrel{\text{def}}{=} X\vec{v}_1 \otimes X\vec{v}_2 \otimes \cdots \otimes X\vec{v}_k.$$

If V affords the character χ and we denote the character of $V^{\otimes k}$ by $\chi^{\otimes k}$, then it is easy to see that $\chi^{\otimes k}(X) = (\chi(X))^k$ for all $X \in GL_n$. In particular, if V is the module for the defining representation then $\chi^{\otimes k}(X) = (x_1 + x_2 + \cdots + x_n)^k$ where the x_i are the eigenvalues of X. Clearly if V corresponds to a polynomial representation then so does $V^{\otimes k}$.

Suppose that V is a module for the defining representation of GL_n. Decomposing $V^{\otimes k}$ into irreducibles produces the following beautiful theorem.

Theorem 7.1.4 ([Scu 01]) *If V is the defining module for GL_n then*

$$V^{\otimes k} \cong \bigoplus_{\substack{\lambda \vdash k \\ l(\lambda) \leq n}} m_\lambda V^\lambda$$

where $m_\lambda = f^\lambda$. $\square$

Taking characters on both sides of Theorem 7.1.4, we immediately obtain

Corollary 7.1.5

$$(x_1 + x_2 + \cdots + x_n)^k = \sum_{\substack{\lambda \vdash k \\ l(\lambda) \leq n}} f^\lambda s_\lambda(\mathbf{x}).\square$$

7.2 Projective characters

The Schur Q-functions give information about the characters of projective representations of S_n. Recall that given a strict partition $\lambda^* \models n$ of length l, there is a single irreducible projective S_n-module S_0^λ when $n - l$ is even and two, S_1^λ and S_{-1}^λ, when $n - l$ is odd. Let ζ_i^λ for $i = 0, \pm 1$ be the corresponding characters. It turns out that these characters are only non-zero on two families of partitions $\mu = (\mu_1, \mu_2, \cdots, \mu_m)$: those where the μ_i are all odd, and those where μ is a strict partition with an odd number of even parts.

The Schur Q-functions only give information about the values of ζ_i^λ on partitions from the first family, but there is an explicit formula, rather than a generating function, for their values on the second (see [Mor 77] for details). If μ has only odd parts, the values of all three characters are the same and we will denote this common value by ζ_μ^λ.

Theorem 7.2.1 ([Scu 11]) *If $\lambda^* \models n$ then*

$$Q_\lambda(\mathbf{x}) = \frac{1}{n!} \sum_{\mu \vdash n} 2^{\lceil \frac{l(\lambda)+l(\mu)}{2} \rceil} c_\mu \, \zeta_\mu^\lambda \, p_\mu(\mathbf{x})$$

where the sum is over all partitions μ with only odd parts, and $\lceil \cdot \rceil$ is the round-up or ceiling function. $\square$

Corollary 7.1.5 also has a projective analog.

Theorem 7.2.2

$$(x_1 + x_2 + \cdots x_n)^k = \sum_{\substack{\lambda^* \models k \\ l(\lambda^*) \leq n}} 2^{-l(\lambda^*)} g^\lambda \, Q_\lambda(\mathbf{x}). \square$$

7.3 Symplectic characters

Let V be a $2n$-dimensional vector space over $\mathbf{C}$ equipped with a non-degenerate skew-symmetric bilinear form $\langle \cdot, \cdot \rangle$. The *symplectic group*, $Sp_{2n} = Sp(V)$, is the subgroup of GL_{2n} that preserves the bilinear form, i.e.,

$$Sp_{2n} = \{X \in GL(V) \mid \langle X\vec{v}, X\vec{w} \rangle = \langle \vec{v}, \vec{w} \rangle \text{ for all } \vec{v}, \vec{w} \in V\}.$$

Polynomial representations and characters are defined as for GL_n. Furthermore, all polynomial representations of Sp_{2n} are completely reducible. Since $X \in Sp_{2n}$ stabilizes a skew-symmetric form, its set of eigenvalues must be of the form $\mathbf{x}^{\pm 1} = \{x_1, {x_1}^{-1}, x_2, {x_2}^{-1}, \cdots, x_n, {x_n}^{-1}\}$. This motivates the symplectic analog of Theorem 7.1.3.

Theorem 7.3.1 *The irreducible polynomial representations of Sp_{2n} are indexed by partitions λ of length at most n. If λ is such a partition with corresponding module $\tilde{V}^\lambda$ then the character afforded by $\tilde{V}^\lambda$ is*

$$\tilde{\phi}^\lambda(X) = sp_\lambda(\mathbf{x}^{\pm 1}),$$

where $\mathbf{x}^{\pm 1} = \{x_1, {x_1}^{-1}, x_2, {x_2}^{-1}, \cdots, x_n, {x_n}^{-1}\}$ is the set of eigenvalues of X. $\square$

Now consider $Sp(V)$ and it's defining module V. We can take tensor powers as before and study the decomposition into irreducibles. In what follows, $\tilde{f}_k^\lambda(n)$ is the number of OYT $(\phi = \lambda^0, \lambda^1, \cdots, \lambda^k)$ such that $l(\lambda^i) \leq n$ for all $i = 1, 2, \cdots, k$.

Theorem 7.3.2 *If V is the defining module for Sp_{2n} then*

$$V^{\otimes k} \cong \bigoplus_{\substack{\lambda \vdash k \\ l(\lambda) \leq n}} m_\lambda \tilde{V}^\lambda$$

where $m_\lambda = \tilde{f}_k^\lambda(n)$. $\square$

Taking characters on both sides above, we obtain:

Corollary 7.3.3

$$(x_1 + {x_1}^{-1} + x_2 + {x_2}^{-1} + \cdots + x_n + {x_n}^{-1})^k = \sum_{\substack{\lambda \vdash k \\ l(\lambda) \leq n}} \tilde{f}_k^\lambda(n) sp_\lambda(\mathbf{x}^{\pm 1}). \ \square$$

Our discussion of symplectic representations has been rather cursory. For more details, see the paper of Sundaram, *Tableaux in representation theory of the classical Lie groups*, elsewhere in this volume.

8 The Knuth correspondence

We will now generalize the Robinson-Schensted map to give combinatorial proofs of Corollary 7.1.5, Theorem 7.2.2, and Corollary 7.3.3. These proofs are due to Knuth [Knu 70], Sagan-Worley [Sag 87, Wor 84], and Berele [Ber 86] respectively.

8.1 Left-justified tableaux

The Robinson-Schensted correspondence is a map between permutations and SYT. To obtain the analog for GYT (which have repeated entries), we will have to introduce permutations with repetitions. A *generalized permutation* is a two line array

$$\pi = \begin{matrix} k_1 & k_2 & \cdots & k_m \\ l_1 & l_2 & \cdots & l_m \end{matrix}$$

which is in lexicographic order where the top line takes precedence. For example,

$$\pi = \begin{matrix} 1 & 1 & 1 & 2 & 2 & 3 \\ 2 & 3 & 3 & 1 & 2 & 1 \end{matrix}.$$

We say that π is a *generalized permutation of* n if every element in π is less than or equal to n. We let GP_n stand for the set of all generalized permutations of n. Associated with each $\pi \in GP_n$ is a pair of contents $(1^{b_1}, 2^{b_2}, \cdots, n^{b_n})$ and $(1^{t_1}, 2^{t_2}, \cdots, n^{t_n})$ where b_i (respectively t_i) is the number of occurences of i in the bottom (respectively top) line of π. The example above has contents $(1^2, 2^2, 3^2)$ and $(1^3, 2^2, 3^1)$. Introducing a new set of variables $\mathbf{y} = \{y_1, y_2, \cdots, y_n\}$ we can define the *monomial of* π to be

$$m(\pi) = {x_1}^{b_1} {x_2}^{b_2} \cdots {x_n}^{b_n} {y_1}^{t_1} {y_2}^{t_2} \cdots {y_n}^{t_n}.$$

Our example has monomial $m(\pi) = {x_1}^2 {x_2}^2 {x_3}^2 {y_1}^3 {y_2}^2 {y_3}^1$.

We claim that the generating function for generalized permutations of n is

$$\sum_{\pi \in GP_n} m(\pi) = \prod_{i,j=1}^{n} \frac{1}{1 - x_i y_j}.$$

To see this, note that

$$\frac{1}{1 - x_i y_j} = 1 + x_i y_j + {x_i}^2 {y_j}^2 + \cdots + {x_i}^k {y_j}^k + \cdots .$$

Thus the term ${x_i}^k {y_j}^k$ corresponds to having the column $\binom{i}{j}$ repeated k times in π.

We should note that there is a bijection between generalized permutations of n and $n \times n$ matrices with non-negative integral entries. This is because having $\binom{i}{j}$ repeated k times in π is equivalent to having the (i, j) entry of the matrix equal to k. Knuth's original proof of Theorem 8.1.1 below was stated in terms of matrices and generalized Young tableaux.

In Schensted's paper [Sch 61], he gave what amounts to a combinatorial proof of Corollary 7.1.5 (although Schur functions were never mentioned explicitly). If remained for Knuth [Knu 70] to give a combinatorial proof of Cauchy's identity, which is a generalization of this corollary, and to make the connection with $s_\lambda(\mathbf{x})$.

Theorem 8.1.1

$$\prod_{i,j=1}^{n} \frac{1}{1 - x_i y_j} = \sum_{l(\lambda) \leq n} s_\lambda(\mathbf{x}) s_\lambda(\mathbf{y}) \ .$$

Proof. We wish to find a bijection

$$\pi \longleftrightarrow (T, U)$$

between generalized permutations $\pi \in GP_n$ and pairs of GYT T, U of the same shape such that the content of T (respectively U) equals the content of the lower (respectively upper) line of π. For the forward direction we form, as before, a sequence of tableaux pairs

$$(T_0, U_0) = (\phi, \phi);\ (T_1, U_1);\ (T_2, U_2);\ \cdots;\ (T_m, U_m) = (T, U)$$

where the elements of the bottom line of π are inserted into the T's and the elements of the top line are placed in the U's. Furtermore the rules of insertion and placement are exactly the same.

Applying this algorithm to the permuation above, we obtain

$$T_i: \quad \phi\ ,\quad 2\ ,\quad 2\,3\ ,\quad 2\,3\,3\ ,\quad \begin{matrix} 1\,3\,3 \\ 2 \\ \end{matrix}\ ,\quad \begin{matrix} 1\,2\,3 \\ 2\,3 \\ \end{matrix}\ ,\quad \begin{matrix} 1\,1\,3 \\ 2\,2 \\ 3 \end{matrix} \quad = T$$

$$U_i: \quad \phi\ ,\quad 1\ ,\quad 1\,1\ ,\quad 1\,1\,1\ ,\quad \begin{matrix} 1\,1\,1 \\ 2 \\ \end{matrix}\ ,\quad \begin{matrix} 1\,1\,1 \\ 2\,2 \\ \end{matrix}\ ,\quad \begin{matrix} 1\,1\,1 \\ 2\,2 \\ 3 \end{matrix} \quad = U.$$

It is easy to verify that the insertion rules make sure that T is a GYT and that U always has weakly increasing rows. To verify the column condition for U, we must make sure that no two equal elements of the top row of π can end up in the same column. But if $k_i = k_{i+1} = k$ in the upper row, then by the lexicographic condition on π we must have $l_i \leq l_{i+1}$. This implies that the insertion path of l_{i+1} will always lie strictly to the right of the path for l_i which gives the desired result. Note that we have shown that all elements equal to k are placed in U from left to right as the algorithm proceeds.

For the inverse map, the only problem is deciding which of the maximum elements of U corresponds to the last insertion. But from the observation just made, the right-most of these maxima is the correct choice to start the deletion process. □

Many of the properties of the Robinson-Schensted algorithm also hold for Knuth's generalization. For details, the reader can consult [Gan 78].

8.2 Shifted tableaux

The analog of Cauchy's identity that corresponds to Theorem 7.2.2 is:

Theorem 8.2.1

$$\prod_{i,j=1}^{n} \frac{1+x_i y_j}{1-x_i y_j} = \sum_{l(\lambda^*)\leq n} 2^{-l(\lambda^*)} Q_\lambda(\mathbf{x}) Q_\lambda(\mathbf{y}) \ .$$

Proof sketch. The left-hand side of this equation counts generalized permutations π where, if a column $\binom{i}{j}$ appears in π, then it's first occurence can be distinguished by being putting a prime on the j. (Picking the '1' or the '$x_i y_i$' in the numerator's $1 + x_i y_j$ corresponds respectively to not distinguishing or distinguishing the j).

To see what the right side counts, remember that primed and unprimed versions of the same integer lie in skew hooks and that the nature of every element of a given hook is completely determined except at it's lower-left end. But $l(\lambda^*)$ skew hooks have their lower-left end on the main diagonal, so $2^{-l(\lambda^*)} Q_\lambda(\mathbf{y})$ counts GST with primes only on off-diagonal elements.

Thus it suffices to find a bijection

$$\pi^* \longleftrightarrow (T^*, U^*)$$

where π^* is a primed generalized permutation of n and T^*, U^* are GST of the same shape such that U only has off-diagonal primes and the content of the lower (respectively upper) line of π^* equals the content of T^* (respectively U^*). Primes are ignored when taking contents. Details of the bijection can be found in [Sag 87] or [Wor 84]. □

8.3 Symplectic tableaux

Berele's algorithm [Ber 86] was constructed to give a combinatorial proof of Corollary 7.3.3, which we restate here.

Theorem 8.3.1

$$(x_1 + {x_1}^{-1} + x_2 + {x_2}^{-1} + \cdots + x_n + {x_n}^{-1})^k = \sum_{\substack{\lambda \vdash k \\ l(\lambda) \leq n}} \tilde{f}_k^\lambda(n) sp_\lambda(\mathbf{x}^{\pm 1}).$$

Proof. The left-hand side counts permutations $\tilde{\pi}$ of length k over the alphabet $\{1 < \bar{1} < 2 < \bar{2} < \cdots < n < \bar{n}\}$ (possibly with repetitions in the lower line, but none in the upper). As expected, $\tilde{\pi}$ has monomial

$$m(\tilde{\pi}) = {x_1}^{m_1} {x_1}^{-\bar{m}_1} {x_2}^{m_2} {x_2}^{-\bar{m}_2} \cdots x_n^{m_n} {x_n}^{-\bar{m}_n}$$

with m_k (respectively $\bar{m}_k$) being the number of k's (respectively $\bar{k}$'s) in $\tilde{\pi}$. Thus we want to give a bijection

$$\tilde{\pi} \longleftrightarrow (\tilde{T}, \tilde{P}_k^\lambda)$$

where $\tilde{\pi}$ is as described above; $\tilde{T}, \tilde{P}_k^\lambda$ are an SPT and an OYT respectively having the same shape λ; and $m(\tilde{\pi}) = m(\tilde{T})$.

In the forward direction, we create a sequence

$$(\tilde{T}_0, \lambda^0) = (\phi, \phi);\ (\tilde{T}_1, \lambda^1);\ (\tilde{T}_2, \lambda^2); \cdots ; (\tilde{T}_k, \lambda^k)$$

so that at the end we can let $\tilde{T} = \tilde{T}_k$ and $\tilde{P}_k^\lambda = (\lambda^0, \lambda^1, \lambda^2, \cdots, \lambda^k)$. We will also build the pairs so that the shapes of $\tilde{T}_i$ and λ^i are the same for all $i = 0, 1, 2, \cdots, k$.

Suppose that $(\tilde{T}_{r-1}, \lambda^{r-1})$ has been constructed and that $\tilde{\pi} = l_1\ l_2\ \cdots\ l_k$. We start row inserting l_r into $\tilde{T}_{r-1}$ as usual. If the symplectic condition (equation (5)) is never violated during the insertion, then we let $\tilde{T}_r = R_{l_r}(\tilde{T}_{r-1})$ and add a box to λ_{r-1} to mark the location of termination.

Suppose, on the other hand, that a violation of equation (5) is about to occur at some point of the insertion. It is easy to see that this could only happen if an i added to row i is trying to bump an $\bar{\imath}$ into row $i+1$. In this case the i and the $\bar{\imath}$ anihalate eachother, creating a (black) hole in cell (i, j) for some j. But now the hole can be filled using an (i, j)-slide, resulting in the new tableau $\tilde{T}_r$ which has one less box. In this case we delete the corresponding box of λ^{r-1} to form λ^r. Note that this cancelation of an i and an $\bar{\imath}$ corresponds to the cancelation of an x_i with an ${x_i}^{-1}$.

By way of illustration, we consider the insertion of $\bar{\imath}$ into the SPT below:

$$\begin{array}{llll} 1 & \bar{1} & 2 & \bar{2} \\ 2 & \bar{2} & 3 & \\ 3 & 3 & \bar{3} & \end{array} \leftarrow \bar{1}, \quad \begin{array}{llll} 1 & \bar{1} & \bar{1} & \bar{2} \\ 2 & \bar{2} & 3 & \\ 3 & 3 & \bar{3} & \end{array} \begin{array}{c} \\ \leftarrow 2 \\ \\ \end{array}, \quad \begin{array}{llll} 1 & \bar{1} & \bar{1} & \bar{2}, \\ 2 & \bullet & 3 & \\ 3 & 3 & \bar{3} & \end{array} \quad \begin{array}{llll} 1 & \bar{1} & \bar{1} & \bar{2}, \\ 2 & 3 & 3 & \\ 3 & \bullet & \bar{3} & \end{array} \quad \begin{array}{llll} 1 & \bar{1} & \bar{1} & \bar{2}\,. \\ 2 & 3 & 3 & \\ 3 & \bar{3} & & \end{array}$$

For a look at the whole algorithm, let's compute the image of $\tilde{\pi} = 2\bar{2}\bar{1}211$:

$$\tilde{T}_r: \quad \phi\,, \quad 2\,, \quad 2\,\bar{2}\,, \quad \begin{array}{l} \bar{1}\,\bar{2} \\ 2 \end{array}, \quad \begin{array}{l} \bar{1}\,2 \\ 2\,\bar{2} \end{array}, \quad \begin{array}{l} 2\,2 \\ \bar{2} \end{array}, \quad 1\,2$$

$$\lambda^r: \quad \phi\,, \quad \square\,, \quad \square\,\square\,, \quad \begin{array}{l} \square\,\square \\ \square \end{array}, \quad \begin{array}{l} \square\,\square \\ \square\,\square \end{array}, \quad \begin{array}{l} \square\,\square \\ \square \end{array}, \quad \square\,\square\,.$$

Thus $\tilde{T} = 1\ 2$ and

$$\tilde{P}_k^\lambda = \quad \phi\,, \quad \square\,, \quad \square\,\square\,, \quad \begin{array}{l} \square\,\square \\ \square \end{array}, \quad \begin{array}{l} \square\,\square \\ \square\,\square \end{array}, \quad \begin{array}{l} \square\,\square \\ \square \end{array}, \quad \square\,\square\,.$$

We leave it as an exercise to the reader to construct the inverse map. $\square$

9 Open questions

Now that the reader has gained some familiarity with tableaux and their relation with representations and symmetric functions, it seems appropriate to propose some outstanding problems using these ideas.

1. Shifted analogs. We have seen (Theorem 4.1.4) that reversing a permutation transposes its P-tableau. What effect does this have on the P^* tableau of the shifted Robinson-Schensted map? More generally, what does it mean to 'transpose' a shifted tableau? There are a host of other problems concerning shifted analogs of known results in the left-justified case. The reader can consult [Sag 87] for further information.

2. Restricted partitions. There are many beautiful product generating functions for various families of tableaux. For example, we have the following theorem of MacMahon:

Theorem 9.0.2 ([MaM 15]) *Fix postitive integers k, l and m and let λ be the rectangular partition (k^l) . Then the generating function for reverse plane partitions with at most k rows, at most l columns (i.e., having shape contained in λ), and with largest part at most m is*

$$\prod_{(i,j)\in\lambda} \frac{1 - x^{h_{i,j}+m}}{1 - x^{h_{i,j}}}\ .\square$$

It would be nice to have a combinatorial proof of this theorem, perhaps using a Hillman-Grassl type bijection. The paper [Sta 86] of Stanley is a good source for problems of this type. His survey article [Sta 71] is also very informative.

3. Projective modules. Since the dimension of the Specht module S^λ is just the number of SYT of shape λ, it is desirable to have a basis constructed out of these tableaux. This can be done using *tabloids* which are equivalence classes of tableaux (two tableaux are equivalent if corresponding rows contain the same set of elements; here, rows and columns need not increase). The symmetric group acts on tabloids in a natural way, and from these permutation modules one can construct the irreducibles. See [Jam 78] for details.

Only recently have the matrices for the irreducible projective representations been constructed by Nazarov [Naz 88]. Can one find a way to use an S_n action on shifted tableaux to accomplish the same task?

4. Hall-Littlewood polynomials. Both the normal Schur functions and the Schur Q-functions are special cases of the Hall-Littlewood polynomials, $Q_\lambda(\mathbf{x};t)$. These polynomials are symmetric in the variables $\mathbf{x}$ with an additional paramenter, t. When $t = 0$ or -1 they specialize to $s_\lambda(\mathbf{x})$ or $Q_\lambda(\mathbf{x})$ respectively. More information about these functions can be found in Macdonald's book [Mad 79] or in the survey article of Morris [Mor 77].

The $Q_\lambda(\mathbf{x};t)$ satisfy the identity

$$\prod_{i,j=1}^{n} \frac{1-tx_iy_j}{1-x_iy_j} = \sum_{l(\lambda^*)\leq n} \frac{1}{b_\lambda(t)} Q_\lambda(\mathbf{x};t)Q_\lambda(\mathbf{y};t)$$

where $b_\lambda(t)$ is a polynomaial in t. This generalizes both Cauchy's identity (Theorem 8.1.1) and Theorem 8.2.1. Perhaps it is possible to find a Robinson-Schensted-Knuth type map to give a combinatorial proof of this result. Thus the left-justified and shifted correspondences would be combined into one.

References

[Ber 86] A. Berele, *A Schensted-type correspondence for the symplectic group*, J. Combin. Theory, Ser. A, 43 (1986), 320-328.

[DKR 78] J. Désarménien, J. P. S. Kung and G.-C. Rota, *Invariant theory, Young bitableaux, and Combinatorics*, Advances in Mathematics, 27 (1978), 63-92.

[DRS 74] P. Doubilet, G.-C. Rota and J. Stein, *On the foundations of combinatorial theory: IX. Combinatorial methods in invariant theory*, Studies in Applied Math., 53 (1974), 185-216.

[FRT 54] J. S. Frame, G. de B. Robinson and R. M. Thrall, *The hooklengths of* S_n, Canad. J. Math., 6 (1954), 316-325.

[F-Z 82] D. S. Franzblau and D. Zeilberger, *A bijective proof of the hook-length formula*, J. of Algorithms, 3 (1982), 317-343.

[Gan 78] E. Gansner, *Matrix correspondences and the enumeration of plane partitions*, Ph.D. thesis, M.I.T., 1978.

[G-M 81] A. Garsia and S. Milne, *A Rogers-Ramanujan bijection*, J. Combin. Theory, Ser. A, 31 (1981), 289-339.

[G-V 85] I. Gessel and G. Viennot, *Binomial determinants, paths, and hooklength formulae*, Adv. in Math., 58 (1985), 300-321.

[G-V ip] I. Gessel and G. Viennot, *Determinants, paths, and plane partitions*, in preparation.

[Gre 74] C. Greene, *An extension of Schensted's theorem*, Advances in Math., 14 (1974), 254-265.

[GNW 79] C. Greene, A. Nijenhuis and H. S. Wilf, *A probabilistic proof of a formula for the number of Young tableaux of a given shape*, Adv. in Math., 31 (1979), 104-109.

[Hai pr] M. D. Haiman, *On mixed insertion, symmetry, and shifted Young tableaux*, preprint.

[H-G 76] A. P. Hillman and R. M. Grassl, *Reverse plane partitions and tableau hook numbers*, J. Combin. Theory, Ser. A, 21 (1976), 216-221.

[Jam 78] G. D. James, *The representation theory of symmetric groups*, Springer-Verlag Lecture Notes in Math., 682 (1978).

[Józ pr] T. Józefiak, *Characters of projective representations of symmetric groups*, preprint.

[Kad pr] K. Kadell, *Schützenberger's "jeu de taquin" and plane partitions*, preprint.

[Knu 70] D. E. Knuth, *Permutations, matrices and generalized Young tableaux*, Pacific J. Math., 34 (1970), 709-727.

[Knu 73] D.E. Knuth, *The art of computer programming, vol. 3: Sorting and searching*, Addison-Wesley Pub., 1973.

[Led 77] W. Ledermann, *Introduction to group characters*, Cambridge University Press, 1977.

[Mad 79] I. G. Macdonald, *Symmetric functions and Hall polynomials*, Oxford University Press, 1979.

[MaL 86] T. J. MacLarnan, *Tableau recursions and symmetric Schensted correspondences for ordinary, shifted and oscillating tableaux*, Ph.D. thesis, UCSD, (1986)

[MaM 15] P. A. MacMahon, *Combinatorial Analysis.* vols. 1 and 2, Cambridge University Press, 1915-1916 (reprinted by Chelsea Pub., 1960).

[Mor 77] A. O. Morris, *A survey of Hall-Littlewood functions and their applications to representation theory*, Combinatoire et Représentation du Groupe Symétrique, D. Foata ed., Springer-Verlag Lecture Notes in Math., 579 (1977), 136-154.

[Naz 88] M. L. Nazarov, *An orthogonal basis of irreducible projective representations of the symmetric group*, Functional Anal. Appl., 22 (1988), 77-78.

[Rem 82] J. B. Remmel, *Bijective proofs of formulae for the number of standard Young tableaux*, Linear and Multilin. Alg., 11 (1982), 45-100.

[Rob 38] G. de B. Robinson, *On the representations of the symmetric group*, Amer. J. Math., 60 (1938), 745-760.

[Sag 79] B. E. Sagan, *An analog of Schensted's algorithm for shifted Young tableaux*, J. Combin. Theory, Ser. A, 27 (1979), 10-18.

[Sag 80] B. E. Sagan, *On selecting a random shifted Young tableau*, J. Algorithms, 1 (1980), 213-234.

[Sag 82] B. E. Sagan, *Enumeration of partitions with hooklengths*, Europ. J. Combinatorics, 3 (1982), 85-94.

[Sag 87] B. E. Sagan, *Shifted tableaux, Schur Q-functions and a conjecture of R. Stanley*, J. Combin. Theory, Ser. A, 45 (1987), 62-103.

[Sag ip] B. E. Sagan, *The symmetric group: representations, combinatorics and symmetric functions*, Wadsworth and Brooks/Cole, in preparation.

[Sch 61] C. Schensted, *Longest increasing and decreasing subsequences*, Canad. J. Math., 13 (1961), 179-191.

[Scu 01] I. Schur, *Über eine Klasse von Matrizen, die sich einer gegeben Matrix zuorden lassen, Dissertation*, Berlin, (1901). Reprinted in Gesammelte Abhandlungen, Springer-Verlag, 1973.

[Scu 11] I. Schur, *Über die Darstellung der symmetrischen und der alterneinden Gruppe durch gebrochene lineare Substitutionen*, J. Reine Angew. Math, 139 (1911), 155-250.

[Scü 63] M. P. Schützenberger, *Quelques remarques sur une construction de Schensted*, Math. Scand., 12 (1963), 117-128.

[Scü 76] M. P. Schützenberger, *La correspondence de Robinson*, Combinatoire et Représentation du Groupe Symétrique, D. Foata ed., Springer-Verlag Lecture Notes in Math., 579 (1977), 59-135.

[Sta 71] R. P. Stanley, *Theory and application of plane partitions I and II*, Stud. in Applied Math., 50 (1971),167-188 and 259-279.

[Sta 72] R. P. Stanley, *Ordered Structures and Partitions*, Memoires of the A.M.S., 119 (1972).

[Sta 86] R. P. Stanley, *A baker's dozen of conjectures on plane partitions*, Combinatoire Enumerative, G. Labelle and P. Leroux eds., Springer-Verlag Lecture Notes in Math., 1234 (1986), 285-296.

[Ste 87] J. R. Stembridge, *Shifted tableaux and the projective representations of symmetric groups*, preprint.

[Sun 86] S. Sundaram, *On the combinatorics of representations of* $Sp(2n, \mathbf{C})$, Ph.D. Thesis, M.I.T., 1978.

[Tho 77] G. P. Thomas, *On a construction of Schützenberger*, Discrete Math. 17 (1977), 107-118.

[Thr 52] R. M. Thrall, *A combinatorial problem*, Michigan Math. J., 1 (1952), 81-88.

[Vie 76] G. Viennot, *Une forme géométrique de la correspondance de Robinson-Schensted*, Combinatoire et Représentation du Groupe Symétrique, D. Foata ed., Springer-Verlag Lecture Notes in Math., 579 (1977), 29-58.

[Vie 84] G. Viennot, *Chain and antichain families grids and Young tableaux*, Annals of Discrete Math., 23 (1984), 409-464.

[Wor 84] D. R. Worley, *A theory of Shifted Young tableaux*, Ph.D. thesis, M.I.T., 1984.

[You 01] A. Young, *Quantitative substitutional analysis I*, Proc. London Math. Soc., 33 (1901), 97-146.

[Zei 84] D. Zeilberger, *A short hook-lengths bijection inspired by the Greene-Nijenhuis-Wilf proof*, Discrete Math., 51 (1984), 101-108.