IV. An Introduction to Symmetric Functions.

THE RING OF SYMMETRIC FUNCTIONS

There is a fair amount of algebraic machinery to be developed, but eventually we will see some quite amazing generating function identities proved by combinatorial methods. After that we'll see an application of these ideas to prove a theorem which, on the face of it, has nothing to do with symmetric functions. (Although there are very strong and important links between symmetric functions and the representation theory of the symmetric and general linear groups, I will say very little about this.)

To begin with, consider the ring $R_n := \mathbb{Z}[x_1, \ldots, x_n]$ of polynomials in n variables with integer coefficients. The symmetric group $\mathcal{S}_n$ acts on $\{x_1, \ldots, x_n\}$ by permutation: for $\sigma \in \mathcal{S}_n$ let $\sigma(x_i) := x_{\sigma(i)}$. This action is extended linearly and multiplicatively to all of R_n: for a polynomial $f(x_1, \ldots, x_n) \in R_n$ we let $(\sigma(f))(x_1, \ldots, x_n) := f(x_{\sigma(1)}, \ldots, x_{\sigma(n)})$. A polynomial $f \in R_n$ is said to be *symmetric* if $\sigma(f) = f$ for all $\sigma \in \mathcal{S}_n$. The set of all symmetric polynomials in R_n is denoted by $\Lambda_n := \mathbb{Z}[x_1, \ldots, x_n]^{\mathcal{S}_n}$.

Exercise: Check that Λ_n is closed under addition and multiplication, and hence forms a subring of R_n, the *ring of symmetric functions in n variables.*

Exercise: Check that the "Rayleigh operator" defined by

$$R(f) := \frac{1}{n!} \sum_{\sigma \in \mathcal{S}_n} \sigma(f)$$

for $f \in R_n$ is a linear transformation $R : R_n \to \Lambda_n \otimes \mathbb{Q}$. (The tensor product with $\mathbb{Q}$ just means that we allow rational coefficients.) Is this a ring homomorphism?

It is profitable to consider all the rings $\{\Lambda_n : n \geq 1\}$ simultaneously as specializations of one "limiting object" Λ, for it is in this limiting case that the proofs become most uniform and special cases are avoided. I shall gloss some details in my treatment – although they are not difficult – as I do not wish to dwell on foundations. For a complete and precise (and concise) treatment, see pp.10-12 of Macdonald. For any $m \geq n$ there is a ring homomorphism $R_m \to R_n$ defined by setting $x_i = 0$ for all $n < i \leq m$; restriction of this to the subring Λ_m defines a ring homomorphism $\rho_{m,n} : \Lambda_m \to \Lambda_n$.

Exercise: Check that if $f \in \Lambda_m$ then $\rho_{m,n}(f) \in \Lambda_n$.

It is clear that if $m \geq n \geq p$ then $\rho_{n,p} \circ \rho_{m,n} = \rho_{m,p}$. Macdonald takes some additional care to show that $\rho_{m,n}$ is always surjective, and moreover if $m \geq n \geq k$ and $\rho_{m,n}$ is restricted to Λ_m^k, the space of symmetric functions in m variables which are homogeneous polynomials of degree k, then $\rho_{m,n}^k : \Lambda_m^k \to \Lambda_n^k$ is **bijective**. (Here, the notation $\rho_{m,n}^k$ just means the restriction of $\rho_{m,n}$ to Λ_m^k.) The symmetric functions which are homogeneous of degree k are formed by taking the the "inverse limit" of $\{\Lambda_n^k : \; n \geq 1\}$ with respect to the homomorphisms $\{\rho^k_{m,n} : \ m \geq n \geq 1\}$. The notation is

$$\Lambda^k := \lim_{\longleftarrow} \Lambda^k_n.$$

What it means is this: an element of Λ^k is an infinite sequence $f = (f_n : \ n \geq 1)$ with $f_n \in \Lambda^k_n$ and such that for all $m \geq n$, $\rho^k_{m,n}(f_m) = f_n$. The *ring of symmetric functions* is the direct sum of these homogeneous pieces:

$$\Lambda := \bigoplus_{k=0}^{\infty} \Lambda^k.$$

Notice that since this is a direct sum, any $f \in \Lambda$ can be written (uniquely) as a finite sum of symmetric functions of homogeneous degree: $f = f_0 + f_1 + \cdots + f_t$ with $f_k \in \Lambda^k$ for each $0 \leq k \leq t$.

Exercise: Check that Λ is a ring.

This is perhaps a rather forbidding definition. In fact, it is best to think of a symmetric function as an infinite sum of monomials in infinitely many indeterminates $\{x_1, x_2, \ldots\}$ which is invariant under all permutations of the variables, and which has an upper bound on the degrees of the monomials which appear in it. As I hope to make clear, in practice the animals which live in Λ form a rich mathematical ecosystem and are not at all the intimidating beasts they might seem to be at first sight.

MONOMIAL SYMMETRIC FUNCTIONS

Our goal in this section is to find a basis for Λ over the integers; that is, a subset $\mathcal{B} \subset \Lambda$ such that every symmetric function can be expressed uniquely as a finite $\mathbb{Z}$-linear combination of elements of $\mathcal{B}$.

Let $\mathbb{P} := \{1, 2, 3, \ldots\}$ and $\mathbb{N} := \{0, 1, 2, \ldots\}$. An *exponent vector* is a function $\alpha : \mathbb{P} \to \mathbb{N}$ such that $|\alpha| := \alpha(1) + \alpha(2) + \cdots$ is finite. The *monomial* associated to an exponent vector α is

$$\mathbf{x}^\alpha := x_1^{\alpha(1)} x_2^{\alpha(2)} \cdots ,$$

which is a finite product. A *partition* is an exponent vector λ which is weakly decreasing: $\lambda(1) \geq \lambda(2) \geq \cdots$. Traditionally, the notation for a partition is

$$\lambda : \ \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k > 0,$$

the λ_i are the *parts* of λ, and k is called the *length* of λ, denoted by $\ell(\lambda)$. The *multiplicity vector* of λ is $\mathbf{m}(\lambda) := \langle m_1, m_2, \ldots \rangle$ in which, for each $j \in \mathbb{P}$, $m_j(\lambda) := \#\{i : \ \lambda_i = j\}$ denotes the number of parts of λ which are equal to j; that is, $m_j(\lambda)$ is the multiplicity of j as a part of λ.

Exercise: For any given $n \in \mathbb{N}$, there are only finitely many partitions λ with $|\lambda| = n$.

In fact, if $p(n)$ denotes the number of partitions with $|\lambda| = n$ then

$$\sum_{n=0}^{\infty} p(n)x^n = \prod_{j=1}^{\infty} \frac{1}{1-x^j}.$$

(This is from C&O 330.)

Let $\mathcal{S}_{\mathbb{P}} := \bigcup_{n=1}^{\infty} \mathcal{S}_n$ be the (little) infinite symmetric group. It is "little" in the sense that the elements of $\mathcal{S}_{\mathbb{P}}$ are all those bijections from $\mathbb{P}$ to $\mathbb{P}$ which leave fixed all but finitely many elements of $\mathbb{P}$. The "big" infinite symmetric group $\mathcal{S}_\infty$ consists of *all* bijections from $\mathbb{P}$ to $\mathbb{P}$. Certainly $\mathcal{S}_{\mathbb{P}}$ is a subgroup of $\mathcal{S}_\infty$.

Exercise: The cardinality of $\mathcal{S}_{\mathbb{P}}$ is $\aleph_0$, while the cardinality of $\mathcal{S}_\infty$ is $2^{\aleph_0}$, so "little" and "big" are appropriate adjectives.

The big infinite symmetric group acts on the set of indeterminates $\{x_1, x_2, \ldots\}$ by permutation: for $\sigma \in \mathcal{S}_\infty$ we have $\sigma(x_i) := x_{\sigma(i)}$ for all $i \in \mathbb{P}$. This extends to an action on the set of monomials:

$$\sigma(\mathbf{x}^\alpha) = \sigma(x_1^{\alpha(1)} x_2^{\alpha(2)} \cdots) := x_{\sigma(1)}^{\alpha(1)} x_{\sigma(2)}^{\alpha(2)} \cdots .$$

Exercise: Notice that the corresponding action of $\mathcal{S}_\infty$ on exponent vectors, such that $\sigma(\mathbf{x}^\alpha) = \mathbf{x}^{\sigma(\alpha)}$, is given by $(\sigma(\alpha))(i) := \alpha(\sigma^{-1}(i))$.

The action of $\mathcal{S}_\infty$ on monomials is extended $\mathbb{Z}$-linearly to all formal $\mathbb{Z}$-linear combinations $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha$ of monomials. Such an expression f is said to be *invariant under a permutaton* $\sigma \in \mathcal{S}_\infty$ provided that $\sigma(f) = f$, and is said to be *symmetric* if it is invariant under all permutations in $\mathcal{S}_\infty$.

Exercise: Let $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha$ be an infinite $\mathbb{Z}$-linear sum of monomials. Show that if f is invariant under all permutations in $\mathcal{S}_{\mathbb{P}}$ then f is invariant under all permutations in $\mathcal{S}_\infty$. (The converse to this exercise is trivial.) This shows that for symmetric functions we can forget about invariance under most of $\mathcal{S}_\infty$ and just use $\mathcal{S}_{\mathbb{P}}$ instead.

For each exponent vector α there is exactly one partition λ such that α may be converted to λ by applying some permutation in $\mathcal{S}_{\mathbb{P}}$. (It is obtained by listing the values of α in weakly decreasing order, each with the appropriate multiplicity.) In other words, the set of all orbits of $\mathcal{S}_{\mathbb{P}}$ acting on the set of exponent vectors (or on the set of monomials) corresponds bijectively with the set of all partitions. The unique partition in the $\mathcal{S}_{\mathbb{P}}$-orbit of α is called the *shape* of α, and is denoted by $\mathrm{sh}(\alpha)$.

Now suppose that $f = \sum_\alpha c_\alpha \mathbf{x}^\alpha$ is a symmetric function in Λ. Since f is a finite sum of its homogeneous pieces $f = f_0 + \cdots + f_t$, we can deal with each homogeneous piece separately; so suppose that f is homogeneous of degree n, say. Let $\mathbf{x}^\alpha$ be a monomial which occurs in f with coefficient $c_\alpha \neq 0$, so $|\alpha| = n$. Since f is invariant under all the permutations in $\mathcal{S}_{\mathbb{P}}$, every monomial $\mathbf{x}^\beta$ in the $\mathcal{S}_{\mathbb{P}}$-orbit of α occurs in f with the same coefficient $c_\beta = c_\alpha$. In particular, if $\lambda = \mathrm{sh}(\alpha)$ then $\mathbf{x}^\lambda$ occurs in f with coefficient $c_\lambda = c_\alpha$.

For any partition λ we define the *monomial symmetric function* m_λ to be

$$m_\lambda := \sum_{\alpha:\ \mathrm{sh}(\alpha)=\lambda} \mathbf{x}^\alpha.$$

That is, m_λ is the sum over all monomials in the same orbit of $\mathfrak{S}_\mathbb{P}$ as $\mathbf{x}^\lambda$. Certainly, $m_\lambda \in \Lambda$. From the above observations, we see that if f is a symmetric function which is homogeneous of degree n, then f may be expressed in the form

$$f = \sum_{\lambda \vdash n} c_\lambda m_\lambda$$

with coefficients $c_\lambda \in \mathbb{Z}$, in which the notation $\lambda \vdash n$ denotes that the summation is over all (finitely many) partitions with $|\lambda| = n$. Moreover, since distinct partitions correspond to distinct orbits of $\mathfrak{S}_\mathbb{P}$ acting on monomials, this expression is unique. Now, since any symmetric function is (uniquely) a finite sum of its homogeneous pieces, we have proved the following.

Proposition 4.1. *The set $\{m_\lambda\}$ (where λ ranges over all partitions) is a $\mathbb{Z}$-basis for Λ. That is, every symmetric function can be written uniquely as a finite $\mathbb{Z}$-linear combination of monomial symmetric functions.*

ELEMENTARY SYMMETRIC FUNCTIONS

Next, we find a set of generators for Λ as a ring, and determine the ring structure of Λ.

For each $j \in \mathbb{N}$, the j-th *elementary symmetric function* e_j is m_{1^j}, where 1^j denotes the partition with j parts all equal to 1. More explicitly,

$$e_j := \sum_{1 \leq i_1 < \cdots < i_j} x_{i_1} \cdots x_{i_j}.$$

In particular, $e_0 = 1$ and $e_1 = x_1 + x_2 + x_3 + \cdots$ and $e_2 = x_1x_2 + x_1x_3 + x_1x_4 + \cdots + x_2x_3 + \cdots$. Another way to view this definition is in terms of the generating function

$$E(t) := \sum_{j=0}^{\infty} e_j t^j = \prod_{i=1}^{\infty} (1 + x_i t).$$

The case $e_0 = 1$ is special; the other e_j are referred to as *one-part* elementary symmetric functions, for reasons which will become clear in a moment. The result we prove in this section is the following.

Proposition 4.2. *The one-part elementary symmetric functions $\{e_1, e_2, \ldots\}$ generate Λ as a ring and are algebraically independent over $\mathbb{Z}$. Therefore,*

$$\Lambda = \mathbb{Z}[e_1, e_2, \ldots]$$

is a polynomial ring over $\mathbb{Z}$ in infinitely many indeterminates $\{e_1, e_2, \ldots\}$.

To say that $\{e_1, e_2, \ldots\}$ generates Λ as a ring over $\mathbb{Z}$ means that every $f \in \Lambda$ can be expressed as a finite $\mathbb{Z}$-linear combination of finite products of the $\{e_1, e_2, \ldots\}$. To say that the one-part elementary symmetric functions are algebraically independent over $\mathbb{Z}$ means that all such expressions are *unique.* Now, if $e_{j_1} e_{j_2} \cdots e_{j_k}$ is a finite product of the $\{e_1, e_2, \ldots\}$, then, since multiplication is commutative, we might as well require that the

FIGURE 1

indices are weakly decreasing. So, for any partition λ we define the *elementary symmetric function* indexed by λ:

$$e_\lambda := e_{\lambda_1} e_{\lambda_2} \cdots e_{\lambda_k}.$$

(Hence the "one-part" terminology above.) Thus, we may restate Proposition 2 as follows:

Proposition 4.3. *The set $\{e_\lambda\}$ (where λ ranges over all partitions) is a $\mathbb{Z}$-basis for Λ. That is, every symmetric function can be written uniquely as a finite $\mathbb{Z}$-linear combination of elementary symmetric functions.*

To prove Proposition 3 we can restrict our attention to Λ^n, the symmetric functions which are homogeneous of degree n, for each $n \in \mathbb{N}$ separately. Since $\{m_\lambda : \lambda \vdash n\}$ is a (finite) basis for Λ^n over $\mathbb{Z}$, we can express each elementary symmetric function of degree n uniquely in terms of this basis. This produces a coefficient matrix which is square of order $p(n)$ (the number of partitions of n). We will show that this matrix is invertible over the integers, and hence is a change-of-basis matrix. Thus, $\{e_\lambda : \lambda \vdash n\}$ will be seen to be a $\mathbb{Z}$-basis for Λ^n, and Proposition 3 follows.

There are a few combinatorial preliminaries to the argument. For any partition λ : $\lambda_1 \geq \cdots \geq \lambda_k > 0$, the *Ferrers diagram* of λ is the set of pairs

$$F_\lambda := \{(a, b) : 1 \leq a \leq \ell(\lambda) \text{ and } 1 \leq b \leq \lambda_a\},$$

visualized graphically as a set of boxes indexed using matrix coordinates (see below). Taking the transpose of F_λ (as if it were a matrix) produces the Ferrers diagram of another partition λ', called the *conjugate* of λ. Certainly, $\lambda'' = \lambda$. Here is an example:

$\lambda :\ 7\ 5\ 4\ 4\ 1$ $\qquad\qquad$ $\lambda' :\ 5\ 4\ 4\ 4\ 2\ 1\ 1$

Exercise: Notice that $m_j(\lambda') = \lambda_j - \lambda_{j+1}$ for all $j \in \mathbb{P}$.

The *reverse lexicographic (revlex) order* on the set of exponent vectors (or on the set of monomials) is defined as follows: $\alpha < \beta$ if either $|\alpha| < |\beta|$, or $|\alpha| = |\beta|$ and at the first index $i \in \mathbb{P}$ for which $\alpha(i) \neq \beta(i)$ we have $\alpha(i) > \beta(i)$. We may restrict revlex order to the set of all partitions, since a partition is a special case of an exponent vector. For

example, for partitions of five we have $5 < 41 < 32 < 311 < 221 < 2111 < 11111$.
Exercise: Notice that among all monomials in the $\mathcal{S}_{\mathbb{P}}$-orbit of $\mathbf{x}^\alpha$, the revlex earliest one is $\mathbf{x}^{\mathrm{sh}(\alpha)}$. Hence, $\mathbf{x}^\lambda$ is the revlex earliest monomial in m_λ.

The following proposition will complete the proof of Proposition 2.

Proposition 4.4. *For any partition* λ,

$$e_{\lambda'} = m_\lambda + \sum_{\lambda < \mu} a_{\lambda\mu} m_\mu,$$

in which the sum is over partitions μ *strictly later than* λ *in revlex order, and the coefficients* $a_{\lambda\mu}$ *are nonnegative integers.*

Before proving this, let's consider as an example the case $n = 3$; the partitions of three are $3 < 21 < 111$ in revlex order. By definition, $e_3 = m_{111}$. To compute $e_{21} = e_2e_1$, notice that we can form any monomial of shape 111 in exactly three ways by taking a term from e_2 times a term from e_1; also, we can form any monomial of shape 21 in exactly one way like this, while partitions of shape 3 can not be formed at all. Hence, $e_{21} = m_{21} + 3m_{111}$. Finally, since $e_{111} = e_1e_1e_1$ we can form any monomial of shape 3 in one way, any monomial of shape 21 in $\binom{3}{2} = 3$ ways, and any monomial of shape 111 in $3! = 6$ ways. Hence, $e_{111} = m_3 + 3m_{21} + 6m_{111}$. Since $(3)' = 111$ and $(21)' = 21$ and $(111)' = 3$, this gives the following equations.

$$\begin{cases} e_{(3)'} & = & m_3 & +3m_{21} & +6m_{111} \\ e_{(21)'} & = & & m_{21} & +3m_{111} \\ e_{(111)'} & = & & & m_{111} \end{cases}$$

In general, with respect to revlex order on the partitions of n, the coefficient matrix $(a_{\lambda\mu})$ in Proposition 4 is upper triangular with ones ($a_{\lambda\lambda} = 1$) on the diagonal. Hence, this matrix is invertible over the integers, and this suffices to prove Proposition 3, and thus Proposition 2.

Proof of Proposition 4. By definition, $e_{\lambda'} = e_{\lambda'_1}e_{\lambda'_2}\cdots e_{\lambda'_t}$ where $t = \lambda_1$ is the length of λ'. This product is clearly a nonnegative integer linear combination of monomials, and so a nonnegative integer linear combination of monomial symmetric functons. Collecting factors with the same index, we see that $e_{\lambda'} = e_1^{m_1}e_2^{m_2}\cdots e_k^{m_k}$ where $m_j := m_j(\lambda')$ is the multiplicity of j as a part of λ'. By an exercise above, $m_j(\lambda') = \lambda_j - \lambda_{j+1}$, and so

$$e_{\lambda'} = e_1^{\lambda_1-\lambda_2}e_2^{\lambda_2-\lambda_3}\cdots e_k^{\lambda_k-\lambda_{k+1}}$$

(where $\lambda_{k+1} = 0$). What is the revlex earliest monomial $\mathbf{x}^\alpha$ which appears in this product? By an exercise above, its exponent vector must be a partition. Thinking about how the revlex order is defined, one sees that the best thing to do is to take the revlex earliest monomial from each factor. That is,

$$x_1^{\lambda_1-\lambda_2}(x_1x_2)^{\lambda_2-\lambda_3}\cdots(x_1x_2\cdots x_k)^{\lambda_k} = x_1^{\lambda_1}x_2^{\lambda_2}\cdots x_k^{\lambda_k} = \mathbf{x}^\lambda.$$

Since the monomial $\mathbf{x}^\lambda$ can be formed in only one way, m_λ occurs in $e_{\lambda'}$ with coefficient one. Since $\mathbf{x}^\lambda$ is the revlex earliest monomial in $e_{\lambda'}$ and m_μ contains the monomial $\mathbf{x}^\mu$, if m_μ occurs in $e_{\lambda'}$ with nonzero coeffient, then $\lambda < \mu$ in revlex order. $\square$

COMPLETE SYMMETRIC FUNCTIONS AND DUALITY

In this section we give another set of generators for the ring Λ, and begin to see some its deeper structure.

For each $j \in \mathbb{N}$, the j-th *complete symmetric function* h_j is defined by

$$h_j := \sum_{\lambda \vdash j} m_\lambda = \sum_{|\alpha|=j} \mathbf{x}^\alpha = \sum_{1 \le i_1 \le \cdots \le i_j} x_{i_1} \cdots x_{i_j}.$$

That is, h_j is the sum over all monomials of degree j. In particular, $h_0 = 1$ and $h_1 = x_1 + x_2 + x_3 + \cdots$ and $h_2 = x_1^2 + x_1x_2 + x_1x_3 + \cdots + x_2^2 + x_2x_3 + \cdots$. Another way to view this definition is in terms of the generating function

$$H(t) := \sum_{j=0}^{\infty} h_j t^j = \prod_{i=1}^{\infty} \left(\frac{1}{1 - x_i t} \right).$$

The case $h_0 = 1$ is special; the other h_j are referred to as *one-part* complete symmetric functions, for reasons which will become clear below.

Since the one-part elementary symmetric functions are algebraically independent over $\mathbb{Z}$, if A is any ring we may define a ring homomorphism $\Lambda \to A$ by choosing the image of e_j in A for each $j \in \mathbb{P}$. (Algebraic independence ensures that any such choices are compatible, and since the $\{e_1, e_2, \ldots\}$ generate Λ, the rest of the homomorphism is determined by $\mathbb{Z}$-linear and multiplicative extension.) Thus, we define a ring homomorphism $\omega : \Lambda \to \Lambda$ by putting $\omega(e_j) := h_j$ for all $j \in \mathbb{P}$.

By the formulae for the generating functions of the one-part elementary or complete symmetric functions, we have $H(t) = E(-t)^{-1}$; this gives $E(-t)H(t) = 1$, or, upon expansion,

$$\sum_{n=0}^{\infty} \left(\sum_{j=0}^{n} (-1)^j e_j h_{n-j} \right) t^n = 1.$$

Proposition 4.5. *The ring homomorphism $\omega : \Lambda \to \Lambda$ defined above is an involution; that is, $\omega^2 = 1$ is the identity function. In particular, ω is a ring automorphism of Λ.*

Proof. To show that $\omega^2 = 1$ it suffices to show that $\omega(h_j) = e_j$ for all $j \in \mathbb{N}$, since the $\{e_1, e_2, \ldots\}$ generate Λ. Since $h_0 = e_0 = 1$, this is trivial for $j = 0$. Now we proceed by induction on $j \in \mathbb{N}$. From $E(-t)H(t) = 1$ we deduce that

$$h_n + \sum_{j=1}^{n} (-1)^j e_j h_{n-j} = 0$$

for all $n \in \mathbb{P}$. Applying ω and using the induction hypothesis, we see that

$$\omega(h_n) + \sum_{j=1}^{n} (-1)^j h_j e_{n-j} = 0.$$

Comparing this with

$$(-1)^n e_n + \sum_{j=1}^{n} (-1)^{n-j} h_j e_{n-j} = 0,$$

we see that $\omega(h_n) = e_n$, completing the proof. □

There is a slick way to do this argument with the generating functions. Since ω is a ring homomorphism and $E(-t)H(t) = 1$, we deduce that $\omega(E(-t))\omega(H(t)) = 1$, by applying ω to the coefficient of each power of t separately. Since $\omega(E(t)) = H(t)$ by definition, this gives $\omega(H(t)) = H(-t)^{-1} = E(t)$ and proves Proposition 5.

Corollary 4.6. *The one-part complete symmetric functions $\{h_1, h_2, \ldots\}$ generate Λ as a ring and are algebraically independent over $\mathbb{Z}$. Therefore,*

$$\Lambda = \mathbb{Z}[h_1, h_2, \ldots]$$

is a polynomial ring over $\mathbb{Z}$ in infinitely many indeterminates $\{h_1, h_2, \ldots\}$.

Now, for any partition λ we define the *complete symmetric function* indexed by λ:

$$h_\lambda := h_{\lambda_1} h_{\lambda_2} \cdots h_{\lambda_k}.$$

(Hence the "one-part" terminology above.) Similarly to the case for elementary symmetric functions, Corollary 6 may be restated as follows.

Corollary 4.7. *The set $\{h_\lambda\}$ (where λ ranges over all partitions) is a $\mathbb{Z}$-basis for Λ. That is, every symmetric function can be written uniquely as a finite $\mathbb{Z}$-linear combination of complete symmetric functions.*

POWER SUM SYMMETRIC FUNCTIONS

We give yet another set of generators; however, at one stage we must introduce denominators in the calculation. Thus, they merely generate the ring $\Lambda \otimes \mathbb{Q}$ of symmetric functions with rational coefficients. The real significance of the power sum symmetric functions is in relation to the character theory of symmetric groups – I will explain this briefly a bit later on.

For each $j \in \mathbb{P}$, define the j-th *one-part power sum symmetric function* to be $p_j := m_j = x_1^j + x_2^j + \cdots$. (Notice that p_j is not defined for $j = 0$.) By now, it is not surprising that for any partition λ we define the *power sum symmetric function* indexed by λ to be

$$p_\lambda := p_{\lambda_1} p_{\lambda_2} \cdots p_{\lambda_k}.$$

The generating function for one-part power sum symmetric functions is

$$\begin{aligned}P(t) \quad &:= \quad \sum_{j=1}^{\infty} p_j t^{j-1} = \sum_{j=1}^{\infty}\sum_{i=1}^{\infty} x_i^j t^{j-1} \\ &= \quad \sum_{i=1}^{\infty} \frac{x_i}{1-x_i t} = \sum_{i=1}^{\infty} \frac{d}{dt} \log \frac{1}{1-x_i t},\end{aligned}$$

from which it follows that

$$P(t) = \frac{d}{dt} \log H(t) = H'(t)/H(t).$$

Exercise: Show that $P(-t) = E'(t)/E(t)$.
From the equation $H'(t) = P(t)H(t)$ we extract the coefficient of t^n, finding that

$$h_n = \frac{1}{n} \sum_{j=0}^{n-1} p_{n-j} h_j$$

for all $n \in \mathbb{P}$. Since $h_0 = 1$, this may also be written as

$$p_n = n h_n - \sum_{j=1}^{n-1} h_{n-j} p_j$$

for all $n \in \mathbb{P}$. By induction on n, these equations imply that for all $n \in \mathbb{P}$,

$$h_n \in \mathbb{Q}[p_1, \ldots, p_n] \quad \text{and} \quad p_n \in \mathbb{Z}[h_1, \ldots, h_n].$$

From this, Proposition 8 and Corollary 9 follow.

Proposition 4.8. *The one-part power sum symmetric functions $\{p_1, p_2, \ldots\}$ generate $\Lambda \otimes \mathbb{Q}$ as a $\mathbb{Q}$-algebra and are algebraically independent over $\mathbb{Q}$. Therefore,*

$$\Lambda \otimes \mathbb{Q} = \mathbb{Q}[p_1, p_2, \ldots]$$

is a polynomial ring over $\mathbb{Q}$ in infinitely many indeterminates $\{p_1, p_2, \ldots\}$.

Corollary 4.9. *The set $\{p_\lambda\}$ (where λ ranges over all partitions) is a $\mathbb{Q}$-basis for $\Lambda \otimes \mathbb{Q}$. That is, every symmetric function with rational coefficients can be written uniquely as a finite $\mathbb{Q}$-linear combination of power sum symmetric functions.*

As an example, consider $h_2 = m_2 + m_{11}$. Since $p_2 = m_2$ and $p_{11} = p_1 p_1 = m_2 + 2m_{11}$, we see that $h_2 = (p_2 + p_1^2)/2$. So although the one-part power sums do generate $\Lambda \otimes \mathbb{Q}$, in which we allow rational coefficients, they do not generate all of Λ if only integer coefficients are allowed.

Exercise: From the identity $\omega(H(t)) = E(t)$ (in which ω is applied to the coefficient of each power of t separately), deduce that for any partition λ of size $|\lambda|$ and length $\ell(\lambda)$,

$$\omega(p_\lambda) = (-1)^{|\lambda| - \ell(\lambda)} p_\lambda.$$

FIGURE 2

Thus the duality map ω has a particularly nice representation with respect to the power sum basis of $\Lambda \otimes \mathbb{Q}$.

SCHUR SYMMETRIC FUNCTIONS

At last, our safari through the jungle of symmetric functions has arrived at the territory of its most prized denizens – the top rung of the evolutionary ladder in this world – the Schur symmetric functions. These correspond to the irreducible representations of the symmetric groups, about which I shall say a few things a bit later on. There are at least three ways to define the Schur functions (and then at least two theorems to prove – that the definitions are equivalent). At this point, we part company from Macdonald's algebraic camp, equip ourselves with our combinatorial gear, and boldly set off in search of our quarry...

Let λ be a partition, and recall that the Ferrers diagram of λ is the set of pairs

$$F_\lambda := \{(a,b) : \ 1 \leq a \leq \ell(\lambda) \text{ and } 1 \leq b \leq \lambda_a\}.$$

A *(Young) tableau of shape* λ is a function $T : F_\lambda \to \mathbb{P}$ satisfying the following two conditions for all $(a,b) \in F_\lambda$:

- if $(a, b+1) \in F_\lambda$ then $T(a,b) \leq T(a,b+1)$;
- if $(a+1, b) \in F_\lambda$ then $T(a,b) < T(a+1,b)$.

We picture the tableau T graphically by writing the value of $T(a,b)$ in the (a,b)-th box of the Ferrers diagram F_λ. Then the first condition says that the entries of the tableau are weakly increasing as we read from left to right in each row, and the second condition says that the entries of the tableau are strictly increasing as we read from top to bottom in each column. Here is an example of a tableau of shape 7 5 4 2 2.

For a tableau T of shape λ, the monomial associated to T is

$$\mathbf{x}^T := \prod_{(a,b) \in F_\lambda} x_{T(a,b)}.$$

(For example, the monomial associated to the tableau above is $x_1^5x_2^3x_3^5x_4x_5^2x_7^2x_8^2$.) The *Schur symmetric function of shape* λ is defined to be

$$s_\lambda := \sum_T \mathbf{x}^T,$$

in which the summation is over all tableau of shape λ. This must strike you as a rather objectionable definition, since it is not at all clear that s_λ is a symmetric function.

Proposition 4.10. *For any partition λ, s_λ is a symmetric function.*

Proof. Certainly, s_λ is a formal sum of monomials which is homogeneous of degree $|\lambda|$, hence is of bounded degree. It remains to show that s_λ is invariant under every permutation in $\mathcal{S}_\mathbb{P}$.
Exercise: Every permutation in $\mathcal{S}_\mathbb{P}$ can be expressed as a composition of a finite sequence of simple transpositions of the form $(i \;\; i+1)$, for various $i \in \mathbb{P}$. (This is certainly not true of all permutations in $\mathcal{S}_\infty$.)
As a consequence of this exercise, to prove the proposition it suffices to show that s_λ is invariant under the simple transposition $(i \;\; i+1)$, for each $i \in \mathbb{P}$.

So, fix any $i \in \mathbb{P}$ and let $\sigma = (i \;\; i+1)$. Let $\mathcal{T}(\lambda)$ be the set of all tableau of shape λ. We are going to define an involution $\psi : \mathcal{T}(\lambda) \to \mathcal{T}(\lambda)$ such that for any tableau T of shape λ, $\mathbf{x}^{\psi(T)} = \sigma(\mathbf{x}^T)$. That is, the monomial associated to $\psi(T)$ is obtained from the monomial associated to T by exchanging the indeterminates x_i and x_{i+1}. Before we define ψ, let's see how this finishes the proof. Since $\psi : \mathcal{T}(\lambda) \to \mathcal{T}(\lambda)$ is a bijection we may argue that

$$s_\lambda = \sum_T \mathbf{x}^T = \sum_T \mathbf{x}^{\psi(T)} = \sum_T \sigma(\mathbf{x}^T) = \sigma(s_\lambda),$$

as required. (Here, the sums are over all tableau T of shape λ.) It remains only to define the involution ψ.

Let T be any tableau of shape λ. Let B be the set of those boxes of F_λ on which T takes the value i or $i+1$. There are at most two boxes of B per column, because the values of T are strictly increasing down each column of F_λ, and if there are two boxes of B in a column then they must be adjacent. Moreover, if both (a, b) and $(a+1, b)$ are in B then $T(a, b) = i$ and $T(a+1, b) = i+1$. Let A be the set of those boxes in B which are not in the same column as any other box of B; this is the set of "adjustable" boxes. Here is an example of such a set B with $i = 6$; the adjustable boxes are left empty.

In each row of T, reading left to right, the entries of the adjustable boxes form a sequence of i-s followed by a sequence of $(i+1)$-s, since the entries of T increase weakly from left to right. For a given row, let's say the adjustable boxes contain p i-s followed by q $(i+1)$-s. Then in $\psi(T)$, the adjustable boxes in this row will contain q i-s followed by p $(i+1)$-s. This transformation is applied simultaneously for every row of T to produce $\psi(T)$.

Now $\psi(T)$ is a tableau of shape λ: certainly the row condition is still fulfilled; also, if (a, b) is an adjustable box then $T(a-1, b) < i$ and $i+1 < T(a+1, b)$ (whenever these boxes

FIGURE 3

are still in F_λ), so $\psi(T)$ satisfies the column condition as well. From the construction it is clear that ψ is an involution. If $x_j \notin \{x_i, x_{i+1}\}$ then the exponent of x_j in $\mathbf{x}^{\psi(T)}$ is the same as the exponent of x_j in $\mathbf{x}^T$, since no values of T other than i and $i+1$ have been changed. After a little thought, one sees that the exponent of x_i in $\mathbf{x}^{\psi(T)}$ equals the exponent of x_{i+1} in $\mathbf{x}^T$, and the exponent of x_{i+1} in $\mathbf{x}^{\psi(T)}$ equals the exponent of x_i in $\mathbf{x}^T$. This shows that $\mathbf{x}^{\psi(T)} = \sigma(\mathbf{x}^T)$, and completes the proof. $\square$

At this stage we are posed with the following challenges. First, since we have determined that the Schur function s_λ really *is* a symmetric function, it has a unique expression in each of the following ways:

- as an integer linear combination of $\{m_\lambda\}$;
- as a polynomial in $\{e_1, e_2, \ldots\}$ with integer coefficients;
- as a polynomial in $\{h_1, h_2, \ldots\}$ with integer coefficients;
- as a polynomial in $\{p_1, p_2, \ldots\}$ with rational coefficients.

What, exactly, are the fomulae in each of these cases? Second, it is natural to wonder whether the set $\{s_\lambda\}$ forms a $\mathbb{Z}$-basis for Λ. Third, we would like to know how the duality map ω acts on s_λ. We'll close this section by knocking off a couple of these problems, saving the others for later sections.

For a tableau T of shape λ, the *content* of T is the exponent vector α of the monomial $\mathbf{x}^T$ associated with T; that is, for each $i \in \mathbb{P}$, $\alpha(i)$ is the number of boxes of F_λ to which T assigns the value i.

Proposition 4.11. *For any partition* λ,

$$s_\lambda = \sum_\mu K_{\lambda\mu} m_\mu,$$

in which $K_{\lambda\mu}$ *is the number of tableau of shape* λ *and content* μ. *In particular, the coefficients* $K_{\lambda\mu}$ *are nonnegative integers.*

Proof. For each partition μ, the coefficient of $\mathbf{x}^\mu$ in s_λ is, by definition, the number of tableau of shape λ and content μ. Since s_λ is a symmetric function, the result follows. $\square$

That was easy! These coefficients $K_{\lambda\mu}$ are known as the *Kostka numbers*; their combinatorial interpretation gives us more information, as follows.

First we define the *dominance (partial) order* on the set of all exponent vectors as follows: $\alpha \succeq \beta$ provided that

$$\alpha(1) + \alpha(2) + \cdots + \alpha(i) \geq \beta(1) + \beta(2) + \cdots + \beta(i)$$

for all $i \in \mathbb{P}$. We may restrict this definition to partitions, as they are a special case of exponent vectors. This is not a total order on exponent vectors of a given weight, not even on partitions of a given weight: *e.g.* the partitions 3 1 1 1 and 2 2 2 of six are mutually incomparable by dominance order.

Exercise: Show that for exponent vectors α, β, if $\alpha \succeq \beta$ in dominance order, then $\alpha \leq \beta$ in revlex order.

Proposition 4.12. *Let* λ *and* μ *be partitions. Then* $K_{\lambda\lambda} = 1$, *and if* $K_{\lambda\mu} \neq 0$ *then* $\lambda \succeq \mu$ *in dominance order.*

Proof. Assume that T is a tableau of shape λ and content μ. Fix any $i \in \mathbb{P}$. There are $\mu_1 + \mu_2 + \cdots + \mu_i$ boxes of F_λ to which T assigns a value in the set $\{1, 2, .., i\}$, by definition of the content of T. These boxes must all occur in the topmost i rows of F_λ, because T is strictly increasing down the columns of F_λ. There are $\lambda_1 + \lambda_2 + \cdots + \lambda_i$ boxes of F_λ in the topmost i rows. Thus, if $K_{\lambda\mu} \neq 0$ then

$$\lambda_1 + \lambda_2 + \cdots + \lambda_i \geq \mu_1 + \mu_2 + \cdots + \mu_i$$

for all $i \in \mathbb{P}$; that is, $\lambda \succeq \mu$. Considering the case in which $\mu = \lambda$, there is a unique tableau of shape λ and content λ – it assigns the value i to every box in the i-th row of F_λ – so $K_{\lambda\lambda} = 1$. $\square$

As an example, consider the case $n = 3$. Using the combinatorial interpretation of the Kostka numbers, it is not difficult to check that

$$\begin{cases} s_3 &= m_3 + m_{21} + m_{111} \\ s_{21} &= \phantom{m_3 +{}} m_{21} + 2m_{111} \\ s_{111} &= \phantom{m_3 + m_{21} + {}} m_{111} \end{cases}$$

Corollary 4.13. *The set* $\{s_\lambda\}$ *(where* λ *ranges over all partitions) is a* $\mathbb{Z}$*-basis for* Λ. *That is, every symmetric function can be written uniquely as a finite* $\mathbb{Z}$*-linear combination of Schur symmetric functions.*

Proof. Fix any $n \in \mathbb{N}$. Let $(K_{\lambda\mu})$ be the square matrix, indexed by partitions of n, giving the coefficients when the Schur functions $\{s_\lambda : \lambda \vdash n\}$ are written as linear combinations of the monomial symmetric functions $\{m_\mu : \mu \vdash n\}$. With the rows and columns listed in revlex order, $(K_{\lambda\mu})$ is upper triangular with ones on the diagonal, hence is invertible over the integers. Thus, for each $n \in \mathbb{N}$, $\{s_\lambda : \lambda \vdash n\}$ is a $\mathbb{Z}$-basis for Λ^n, proving the result. □

In fact, $\{s_\lambda\}$ is the nicest basis for Λ that there is, and it is essentially unique. We'll have to wait a little while longer before this can be fully explained.

JACOBI-TRUDY AND GESSEL-VIENNOT

In this section we express the Schur functions as polynomials in the one-part elementaries, and as polynomials in the one-part completes. This allows us to determine how the duality map acts on Schur functions. The derivations rely on interpreting the one-part elementaries and completes as generating functions for lattice paths in the integer plane.

Consider $\mathbb{Z}^2$ as a directed graph in which the vertices are the ordered pairs $(a, b) \in \mathbb{Z}^2$ and there are directed *horizontal edges* $(a, b) \to (a + 1, b)$ and *vertical edges* $(a, b) \to (a, b+1)$ for all $(a, b) \in \mathbb{Z}^2$. For us, a *lattice path* is an infinite path in $\mathbb{Z}^2$ which has only finitely many horizontal edges. If the lattice path P begins at (a, b) and has j horizontal edges, we say it is a path *from* (a, b) *to* $(a + j, \infty)$.

We define two subgraphs of $\mathbb{Z}^2$ with labels on the horizontal edges, as follows. The subgraph $\mathcal{E}$ of $\mathbb{Z}^2$ has as vertices those pairs (a, b) with $a + b \geq 0$, and the horizontal edge $(a, b) \to (a + 1, b)$ is given the label $a + b + 1$. We call $\mathcal{E}$ the *elementary grid.* Notice that the translation $(a, b) \mapsto (a+1, b-1)$ generates an infinite cyclic group of automorphisms of $\mathcal{E}$. The subgraph $\mathcal{H}$ of $\mathbb{Z}^2$ has as vertices those pairs (a, b) with $b \geq 0$, and the horizontal edge $(a, b) \to (a + 1, b)$ is given the label $b + 1$. We call $\mathcal{H}$ the *complete grid.* Notice that the translation $(a, b) \mapsto (a + 1, b)$ generates an infinite cyclic group of automorphisms of $\mathcal{H}$. Patches of these two grids are shown below.

For either grid $\mathcal{E}$ or $\mathcal{H}$, think of each horizontal edge labelled $i \in \mathbb{P}$ as having a copy of the indeterminate x_i on it. For a lattice path P in either grid, the *monomial associated to* P is the product of the indeterminates on its horizontal edges; this will be denoted either $\mathbf{x}^{\mathcal{E}(P)}$ or $\mathbf{x}^{\mathcal{H}(P)}$, depending on the grid. Proposition 14 is immediate from the definitions.

Proposition 4.14. *Let* $j \in \mathbb{N}$.
(a) *We have* $e_j = \sum_P \mathbf{x}^{\mathcal{E}(P)}$*, summing over all lattice paths from* $(0, 0)$ *to* (j, ∞).
(b) *We have* $h_j = \sum_P \mathbf{x}^{\mathcal{H}(P)}$*, summing over all lattice paths from* $(0, 0)$ *to* (j, ∞).

Proof. In the grid $\mathcal{E}$ there is an obvious bijection between j-element subsets of $\mathbb{P}$ and lattice paths from $(0, 0)$ to (j, ∞). In the grid $\mathcal{H}$ there is an obvious bijection between j-element multisubsets of $\mathbb{P}$ and lattice paths from $(0, 0)$ to (j, ∞). □

Now consider the Schur function $s_\lambda = \sum_T \mathbf{x}^T$ indexed by the partition λ. Let T be any tableau of shape λ. For each $1 \leq i \leq \ell(\lambda)$, since the i-th row of T is weakly increasing

Figure 4

it determines a lattice path P_i in $\mathcal{H}$ from $(a_i, 0)$ to $(a_i + \lambda_i, \infty)$; here we can choose any $a_i \in \mathbb{Z}$ because of the automorphisms of $\mathcal{H}$. We'll choose the a_i sensibly, so that the column condition on T is translated into a combinatorial condition on the sequence of lattice paths $(P_1, \ldots, P_{\ell(\lambda)})$.

Without loss of generality we can start P_1 at $(-1, 0)$. Now, because of the column condition on T, the i-th horizontal step of P_2 must have a label strictly larger than the label of the i-th horizontal step of P_1. This can be ensured if we begin P_2 at the point $(-2, 0)$ and require that P_1 and P_2 must be vertex-disjoint. Similarly, if we begin P_3 at $(-3, 0)$ and require that P_2 and P_3 are vertex-disjoint, then the column condition for T will be satisfied between rows two and three.

Proposition 4.15. *Let λ be a partition of length k. There is a bijection between the set of tableau T of shape λ and the set of k-tuples of lattice paths $(P_1, \ldots, P_k)$ in which*
- *each P_i goes from $(-i, 0)$ to $(\lambda_i - i, \infty)$,*
- *these lattice paths are pairwise vertex-disjoint.*

Moreover, if T corresponds to $(P_1, \ldots, P_k)$ in this bijection, then

$$\mathbf{x}^T = \mathbf{x}^{\mathcal{H}(P_1)} \mathbf{x}^{\mathcal{H}(P_2)} \cdots \mathbf{x}^{\mathcal{H}(P_k)}.$$

Proof. This is left as an exercise. □

As an example, here is a picture of the 5-tuple of lattice paths in $\mathcal{H}$ corresponding to the tableau on page 10.

FIGURE 5

Thus we interpret s_λ as the sum over $\ell(\lambda)$-tuples of pairwise nonintersecting lattice paths in $\mathcal{H}$, as in the statement of Proposition 15. This is getting us closer to expressing s_λ as a polynomial in $\{h_1, h_2, \ldots\}$, but what next?

The last step in the argument is the following beautiful theorem of Gessel and Viennot. (They did not state it this generally, but the proof below is essentially theirs.) Begin with a finite, acyclic, directed graph $\mathcal{G}$. Assume that each edge e of $\mathcal{G}$ has associated with it a weight x_e in some commutative ring R. For any path P in $\mathcal{G}$, let $\mathbf{x}^P$ be the product of the weights of the edges in P. For any two vertices (v, w) of $\mathcal{G}$, let $M(v, w)$ be the sum of $\mathbf{x}^P$ over all paths P in $\mathcal{G}$ from v to w. (Notice that the hypotheses on $\mathcal{G}$ ensure that $M(v, w)$ is a polynomial in the edge-weights.)

Proposition 4.16 (Gessel-Viennot)**.** *Let $\mathcal{G}$ be a finite, acyclic, directed graph which is properly embedded in the plane, with edge-weights $\{x_e\}$ in a commutative ring R. Let $A_1, \ldots, A_k, Z_k, \ldots, Z_1$ be pairwise distinct vertices on the outer face of $\mathcal{G}$ which appear in that cyclic order around the outer face. Let*

$$G(A_1, \ldots, A_k; Z_1, \ldots, Z_k) := \sum_{(P_1,\ldots,P_k)} \mathbf{x}^{P_1}\mathbf{x}^{P_2}\cdots\mathbf{x}^{P_k},$$

in which the sum is over all k-tuples of paths in $\mathcal{G}$ such that
- *P_i goes from A_i to Z_i for each $1 \leq i \leq k$, and*
- *these paths are pairwise vertex-disjoint.*

Then

$$G(A_1, \ldots, A_k; Z_1, \ldots, Z_k) = \det(M(A_i, Z_j)).$$

(On the right side, $(M(A_i, Z_j))$ is a k-by-k matrix.)

Proof. Expanding the determinant on the right side, we have

$$\begin{aligned} \det(M(A_i, Z_j)) &= \sum_{\sigma \in \mathcal{S}_k} \mathrm{sgn}(\sigma) \prod_{i=1}^{k} M(A_i, Z_{\sigma(i)}) \\ &= \sum_{(P_1, \ldots, P_k)} \mathrm{sgn}(P_1, \ldots, P_k) \mathbf{x}^{P_1} \mathbf{x}^{P_2} \cdots \mathbf{x}^{P_k}, \end{aligned}$$

in which the last sum is over all k-tuples of paths in $\mathcal{G}$ such that for some permutation $\sigma \in \mathcal{S}_k$, P_i goes from A_i to $Z_{\sigma(i)}$ for each $1 \leq i \leq k$; in this expression we put $\mathrm{sgn}(P_1, \ldots, P_k) := \mathrm{sgn}(\sigma)$ for the unique such permutation. We will show that the contribution of the k-tuples of paths which are not vertex-disjoint sums to zero.

Let $\mathcal{P}$ be the set of all k-tuples as in the previous paragraph. We will define an involution $\psi : \mathcal{P} \to \mathcal{P}$ such that if $\psi(P_1, \ldots, P_k) = (Q_1, \ldots, Q_k)$ then either
- $(Q_1, \ldots, Q_k) = (P_1, \ldots, P_k)$, or
- $\mathrm{sgn}(Q_1, \ldots, Q_k) = -\mathrm{sgn}(P_1, \ldots, P_k)$ and $\mathbf{x}^{Q_1} \mathbf{x}^{Q_2} \cdots \mathbf{x}^{Q_k} = \mathbf{x}^{P_1} \mathbf{x}^{P_2} \cdots \mathbf{x}^{P_k}$.

From this it will follow that in the expansion of the determinant $\det(M(A_i, Z_j))$ above, we need only sum over those k-tuples in $\mathcal{P}$ which are left fixed by ψ.

Here is the definition of $\psi : \mathcal{P} \to \mathcal{P}$. Consider any k-tuple $(P_1, \ldots, P_k) \in \mathcal{P}$ with $\sigma \in \mathcal{S}_k$ such that P_i goes from A_i to $Z_{\sigma(i)}$ for each $1 \leq i \leq k$. If this k-tuple of paths is pairwise vertex-disjoint then it is left fixed by ψ. Otherwise, let $1 \leq t \leq k$ be the first index for which P_t shares a vertex with some other path in the k-tuple. Traversing P_t from A_t towards $Z_{\sigma(t)}$, let V be the first vertex of P_t which is on some other path in the k-tuple. Let $t < u \leq k$ be the first index such that V is on the path P_u. Now we define $\psi(P_1, \ldots, P_k) = (Q_1, \ldots, Q_k)$ as follows. If $i \notin \{t, u\}$ then $Q_i := P_i$. We let Q_t consist of the segment of P_t from A_t to V, followed by the segment of P_u from V to $Z_{\sigma(u)}$; similarly, we let Q_u consist of the segment of P_u from A_u to V, followed by the segment of P_t from V to $Z_{\sigma(t)}$.

Exercise: Check that $\psi : \mathcal{P} \to \mathcal{P}$ is an involution.

Exercise: Check the properties claimed for ψ in the previous paragraph.

Here is a schematic picture for the construction of ψ.

After the cancellation of terms according to ψ, a k-tuple in $\mathcal{P}$ contributes to the expansion of $\det(M(A_i, Z_j))$ if and only if it is pairwise vertex-disjoint. Since $\mathcal{G}$ is properly embedded in the plane and $A_1, \ldots, A_k, Z_k, \ldots, Z_1$ are arranged cyclically on the outer face of $\mathcal{G}$, this implies that P_i goes from A_i to Z_i for all $1 \leq i \leq k$. Hence the sign of each of these terms is positive, and the result follows. $\square$

FIGURE 6

It remains only to apply the Gessel-Viennot Theorem on the complete grid $\mathcal{H}$. The result is known as the Jacobi-Trudy Formula.

Proposition 4.17. *Let λ be a partition of length k. Then $s_\lambda = \det(h_{\lambda_i - i + j})$ in which the right side is a k-by-k determinant and, by convention, $h_r = 0$ if $r < 0$.*

Proof. From Proposition 15, s_λ is the generating function for k-tuples $(P_1, \ldots, P_k)$ of pairwise vertex-disjoint lattice paths P_i from $(-i, 0)$ to $(\lambda_i - i, \infty)$ in $\mathcal{H}$. Now fix any (large) $N \in \mathbb{N}$, and specialize $x_i = 0$ for all $i > N$. Thus $s_\lambda(x_1, \ldots, x_N)$ is a symmetric function in finitely many variables. In the lattice path picture, if any of the P_i uses an edge with label $i > N$ then the contribution of $(P_1, \ldots, P_k)$ to s_λ is zero. So after this specialization, s_λ is the generating function for k-tuples $(P_1, \ldots, P_k)$ of pairwise vertex-disjoint lattice paths P_i from $(-i, 0)$ to $(\lambda_i - i, N)$ in the finite subgraph $\mathcal{H}(\lambda, N)$ of $\mathcal{H}$ induced by those vertices (a, b) such that $-k \leq a \leq \lambda_1 - 1$ and $0 \leq b \leq N$. One checks that the hypotheses of the Gessel-Viennot Theorem are satisfied with $\mathcal{G} = \mathcal{H}(\lambda, N)$, $A_i = (-i, 0)$, and $Z_j = \lambda_j - j$, and that $M(A_i, Z_j) = h_{\lambda_j - j + i}(x_1, \ldots, x_N)$ in this case. Therefore

$$s_\lambda(x_1, \ldots, x_N) = \det(h_{\lambda_j - j + i}(x_1, \ldots, x_N)).$$

To complete the proof, just transpose the matrix on the right side and take the limit as N tends to infinity. □

For example, for the partition 5 3 2 2,

$$s_{5322} = \det \begin{bmatrix} h_5 & h_6 & h_7 & h_8 \\ h_2 & h_3 & h_4 & h_5 \\ 1 & h_1 & h_2 & h_3 \\ 0 & 1 & h_1 & h_2 \end{bmatrix}.$$

Analogously, there is a dual form of the Jacobi-Trudi identity expressing s_λ as a polynomial in the one-part complete symmetric functions.

Proposition 4.18. *Let λ be a partition. Then $s_\lambda = \det(e_{\lambda'_i - i + j})$ in which the right side is a λ_1-by-λ_1 determinant and, by convention, $e_r = 0$ if $r < 0$.*

Proof. This is an exercise on homework assignment 3. □

For example, for the partition 5 3 2 2, since the conjugate partition is 4 4 2 1 1,

$$s_{5322} = \det \begin{bmatrix} e_4 & e_5 & e_6 & e_7 & e_8 \\ e_3 & e_4 & e_5 & e_6 & e_7 \\ 1 & e_1 & e_2 & e_3 & e_4 \\ 0 & 0 & 1 & e_1 & e_2 \\ 0 & 0 & 0 & 1 & e_1 \end{bmatrix}.$$

Corollary 4.19. *For any partition, $\omega(s_\lambda) = s_{\lambda'}$.*

Proof. By Jacobi-Trudy and its dual form,

$$\omega(s_\lambda) = \omega(\det(e_{\lambda'_i - i + j})) = \det(h_{\lambda'_i - i + j}) = s_{\lambda'}.$$

□

THE CAUCHY PRODUCT AND ORTHOGONALITY

In this section we define a symmetric, positive definite, integer-valued, bilinear form on the ring Λ of symmetric functions.

We will be working with symmetric functions in two sets of indeterminates, $\mathbf{x} = \{x_1, x_2, \ldots\}$ and $\mathbf{y} = \{y_1, y_2, \ldots\}$, so we will use notation such as $h_\lambda(\mathbf{x})$, $m_\lambda(\mathbf{y})$, $p_\lambda(\mathbf{x})$, *etc.* to indicate which indeterminates are being considered. Similarly, we will specify the set of indeterminates with the notations $H(\mathbf{x}; t)$, $P(\mathbf{y}; t)$, *etc.* for the generating functions for one-part symmetric functions.

To begin, we give two expansions in symmetric functions of the *Cauchy product*

$$\Pi(\mathbf{x}, \mathbf{y}) := \prod_{i=1}^{\infty} \prod_{j=1}^{\infty} \frac{1}{1 - x_i y_j}.$$

Proposition 4.20. *The Cauchy product can be expanded as follows.*
(a) $\Pi(\mathbf{x}, \mathbf{y}) = \sum_\lambda h_\lambda(\mathbf{x}) m_\lambda(\mathbf{y}) = \sum_\lambda m_\lambda(\mathbf{x}) h_\lambda(\mathbf{y})$.
(b) $\Pi(\mathbf{x}, \mathbf{y}) = \sum_\lambda p_\lambda(\mathbf{x}) p_\lambda(\mathbf{y}) / z_\lambda$.
In these expressions, the sums are over all partitions, and for a partition λ with multiplicity vector $\mathbf{m}(\lambda) = \langle m_1, m_2, \ldots \rangle$,

$$z_\lambda := \prod_{j=1}^{\infty} m_j! j^{m_j}.$$

Proof. For part (a) we calculate as follows.

$$\begin{aligned} \Pi(\mathbf{x}, \mathbf{y}) &= \prod_{j=1}^{\infty} H(\mathbf{x}; y_j) = \prod_{j=1}^{\infty} \left(\sum_{r=0}^{\infty} h_r(\mathbf{x}) y_j^r \right) \\ &= \sum_{\alpha} h_\alpha(\mathbf{x}) \mathbf{y}^\alpha = \sum_{\lambda} h_\lambda(\mathbf{x}) m_\lambda(\mathbf{y}). \end{aligned}$$

In the next-to-last summation, α ranges over all multiplicity vectors and $h_\alpha := \prod_i h_{\alpha(i)}$. The other expression follows by symmetry.

For part (b), recall that $P(t) = (d/dt) \log H(t)$. (The set of indeterminates does not matter yet.) Integrating with respect to t we get

$$\sum_{j=1}^{\infty} \frac{p_j t^j}{j} = \log H(t).$$

(Notice that both sides have zero constant term, so the constant of integration has been chosen correctly.) Exponentiating, we find that

$$\begin{aligned} H(t) &= \exp\left(\sum_{j=1}^{\infty} \frac{p_j t^j}{j}\right) = \prod_{j=1}^{\infty} \exp\left(\frac{p_j t^j}{j}\right) \\ &= \prod_{j=1}^{\infty} \left(\sum_{m=0}^{\infty} \frac{(p_j t^j)^m}{m! j^m}\right) = \sum_{\lambda} \frac{p_\lambda t^{|\lambda|}}{z_\lambda}. \end{aligned}$$

Now, letting $\mathbf{xy}$ denote the set $\{x_i y_j\}$ we see that $\Pi(\mathbf{x}, \mathbf{y}) = H(\mathbf{xy}; 1)$. Also, for any $r \in \mathbb{P}$,

$$p_r(\mathbf{xy}) = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} (x_i y_j)^r = p_r(\mathbf{x}) p_r(\mathbf{y}),$$

from which it follows that $p_\lambda(\mathbf{xy}) = p_\lambda(\mathbf{x}) p_\lambda(\mathbf{y})$ for any partition λ. From this, part (b) follows. $\square$

Now we define a bilinear form $\langle \cdot, \cdot \rangle$ on Λ by requiring that $\{h_\lambda\}$ and $\{m_\lambda\}$ be dual bases; that is, for any partitions λ, μ:

$$\langle h_\lambda, m_\mu \rangle := \delta_{\lambda\mu},$$

in which $\delta_{\lambda\mu}$ is the Kronecker delta function. From this definition it is clear that $\langle \cdot, \cdot \rangle$ is bilinear, integer-valued, and that in the direct sum $\Lambda = \bigoplus_{n=0}^{\infty} \Lambda^n$ of Λ as a sum of its homogeneous pieces, the homogeneous pieces Λ^n are pairwise orthogonal. Symmetry and positive definiteness of $\langle \cdot, \cdot \rangle$ are a consequence of the following rather surprising fact.

Proposition 4.21. *Let $\{v_\lambda\}$ and $\{w_\lambda\}$ be $\mathbb{Q}$-bases of $\Lambda \otimes \mathbb{Q}$, such that each v_λ and w_λ is homogeneous of degree $|\lambda|$. The following conditions are equivalent:*
(a) *The bases $\{v_\lambda\}$ and $\{w_\lambda\}$ are dual with respect to $\langle \cdot, \cdot \rangle$; that is, $\langle v_\lambda, w_\mu \rangle = \delta_{\lambda\mu}$ for all partitions λ, μ;*
(b) $\Pi(\mathbf{x}, \mathbf{y}) = \sum_\lambda v_\lambda(\mathbf{x}) w_\lambda(\mathbf{y})$

Proof. Certainly $\langle v_\lambda, w_\mu \rangle = 0$ if $|\lambda| \neq |\mu|$. Fix any $n \in \mathbb{N}$ and let $\lambda, \mu \vdash n$. Express v_λ in terms of the basis $\{h_\nu\}$ and w_μ in terms of the basis $\{m_\theta\}$; we get square matrices $A = (a_{\lambda\nu})$ and $B = (b_{\mu\theta})$ such that

$$v_\lambda = \sum_\nu a_{\lambda\nu} h_\nu \quad \text{and} \quad w_\mu = \sum_\theta b_{\mu\theta} m_\theta.$$

By definition of $\langle \cdot, \cdot \rangle$,

$$\langle v_\lambda, w_\mu \rangle = \sum_{\nu,\theta} a_{\lambda\nu} b_{\mu\theta} \langle h_\nu, m_\theta \rangle = \sum_\nu a_{\lambda\nu} b_{\mu\nu} = (AB^\dagger)_{\lambda\mu}.$$

(Here $B^\dagger$ denotes the transpose of B.) So condition (a) is equivalent to the statement that for all $n \in \mathbb{N}$, $AB^\dagger = I$, the identity matrix. Also,

$$\begin{aligned}\sum_\lambda v_\lambda(\mathbf{x})w_\lambda(\mathbf{y}) &= \sum_\lambda \left(\sum_\nu a_{\lambda\nu}h_\nu(\mathbf{x})\right)\left(\sum_\theta b_{\lambda\theta}m_\theta(\mathbf{y})\right) \\ &= \sum_{\nu,\theta}\left(\sum_\lambda a_{\lambda\nu}b_{\lambda\theta}\right)h_\nu(\mathbf{x})m_\theta(\mathbf{y}).\end{aligned}$$

Thus, condition (b) is equivalent to the statement that for all $n \in \mathbb{N}$, $A^\dagger B = I$. Therefore conditions (a) and (b) are equivalent. □

Corollary 4.22. *The bilinear form $\langle\cdot,\cdot\rangle$ is symmetric and positive definite, and ω is an isometry.*

Proof. From Proposition 20(b) and Proposition 21, the basis $\{p_\lambda\}$ of $\Lambda\otimes\mathbb{Q}$ is orthogonal:

$$\langle p_\lambda, p_\mu\rangle = z_\lambda\delta_{\lambda\mu}.$$

Since $z_\lambda > 0$ for all λ, it follows that the bilinear form is symmetric and positive definite. Since $\omega(p_\lambda) = (-1)^{|\lambda|-\ell(\lambda)}$ it is easy to check that for any partitions λ, μ:

$$\langle\omega(p_\lambda), \omega(p_\mu)\rangle = \langle p_\lambda, p_\mu\rangle.$$

Hence, ω is an isometry. □

THE ROBINSON-SCHENSTED CORRESPONDENCE

In this section we prove combinatorially that the Schur functions form an orthonormal basis of the ring of symmetric functions. By the results of the previous section, the two conditions in Proposition 23 are equivalent.

Proposition 4.23. (a) $\Pi(\mathbf{x},\mathbf{y}) = \sum_\lambda s_\lambda(\mathbf{x})s_\lambda(\mathbf{y})$.
(b) *The Schur functions $\{s_\lambda\}$ form an orthonormal $\mathbb{Z}$-basis for Λ.*

Before proving this, notice that any other orthonormal $\mathbb{Z}$-basis of Λ is related to $\{s_\lambda\}$ by (for each $n \in \mathbb{N}$) an orthogonal matrix with integer coefficients.
Exercise: Prove that if M is a matrix of integers and $MM^\dagger = I$ then M is a signed permutation matrix.
So every orthonormal basis of Λ is obtained from the Schur function basis by possibly reindexing and changing the signs of some of the $\{s_\lambda\}$. This is the sense in which I meant that the Schur function basis is essentially unique. Since each s_λ is constructed naturally from the indexing partition λ, the choice of signs and indices of $\{s_\lambda\}$ is a sensible one.

We will prove Proposition 23(a) combinatorially by establishing a bijection between the objects counted by the generating functions on each side of the equation. On the right side, $\sum_\lambda s_\lambda(\mathbf{x})s_\lambda(\mathbf{y})$ is the generating function for all pairs (T, U) of tableaux in which T and U have the same shape; the pair (T, U) contributes $\mathbf{x}^T\mathbf{y}^U$ to the right side. On the left side, $\Pi(\mathbf{x},\mathbf{y}) = \prod_{i,j}(1 - x_i y_j)^{-1}$ is the generating function for all finite sequences of pairs $\beta = ((i_1, j_1), \ldots, (i_n, j_n))$ such that $j_1 \leq j_2 \leq \cdots \leq j_n$ and if $j_t = j_{t+1}$ then $i_t \leq i_{t+1}$; for convenience we call such a sequence β a *biword.* The biword β contributes $(\mathbf{xy})^\beta := \prod_{t=1}^n x_{i_t} y_{j_t}$ to the left side. To prove Proposition 23(a) we establish a bijection between biwords β and pairs of tableaux (T, U) of the same shape, such that if β corresponds to (T, U) then $(\mathbf{xy})^\beta = \mathbf{x}^T \mathbf{y}^U$. This is the Robinson-Schensted correspondence.

The construction is admittedly a bit baroque, but at heart it is simply founded on extensive iteration of one basic operation, called row insertion. Let $R: \ r_1 \leq r_2 \leq \cdots \leq r_k$ and t be positive integers. (Think of R as being one row in a tableau). The *insertion of t into row R* is defined as follows. If $r_k \leq t$ then we just append t to row R to get $R': \ \ r_1 \leq r_2 \leq \cdots \leq r_k \leq t$; also, we will define $t' := 0$ in this case. Otherwise, $t < r_k$, so let $1 \leq u \leq k$ be the greatest index for which $r_{u-1} \leq t$ (with the convention that $r_0 = 1$) In this case we put $R': \ r_1 \leq \cdots \leq r_{u-1} \leq t < r_{u+1} \leq \cdots \leq r_k$ and let $t' := r_u$. We use the notation $(t' |\, R') \Longleftarrow (R |\, t)$ for the operation of row insertion. For example,

$$.(7|\, 2,3,5,5,6,8,9) \Longleftarrow (2,3,5,5,7,8,9|\, 6)$$

A nice way to visualize this is to let t sweep across R from right to left, looking for the first place it will fit, and displacing the element of R which was previously in that position.
Exercise: Show that if $(t' |\, R') \Longleftarrow (R |\, t)$ then either $t' = 0$ or $t' > t$.

We are going to iterate row insertion by applying it successively to each row in a tableau T. For ease of expression, the convention $(0|\, R) \Longleftarrow (R|\, 0)$ for any row R is useful. So, let T be a tableau of shape λ, and let the rows of T be $R_1, \ldots, R_k$, in which $k = \ell(\lambda)$. Let $t = t_1$ be a positive integer. Perform the row insertions

$$\begin{array}{rcl}
(t_2|\, R_1') & \Longleftarrow & (R_1|\, t_1) \\
(t_3|\, R_2') & \Longleftarrow & (R_2|\, t_2) \\
& \vdots & \\
(t_{k+1}|\, R_k') & \Longleftarrow & (R_k|\, t_k) \\
.(0|\, R_{k+1}') & \Longleftarrow & (|\, t_{k+1})
\end{array}$$

We obtain a sequence of rows $(R_1', \ldots, R_{k+1}')$, including R_{k+1}' only if it is not empty. We will show that these are the rows of a tableau T'; this defines the *insertion of t into the tableau T*, which we denote by $(T \leftarrow t)$. For example, consider the tableau of shape 5 3 2 2 with rows $(1, 1, 2, 4, 5), (2, 4, 5), (3, 5), (6, 7)$, and insert the positive integer 3:

$$\begin{array}{rcl} (4|\,1,1,2,3,5) & \Longleftarrow & (1,1,2,4,5|\,3) \\ (5|\,2,4,4) & \Longleftarrow & (2,4,5|\,4) \\ (0|\,3,5,5) & \Longleftarrow & (3,5|\,5) \\ (0|\,6,7) & \Longleftarrow & (6,7|\,0) \\ .(0|\,) & \Longleftarrow & (|\,0) \end{array}$$

The result is the tableau T' of shape 5 3 3 2 with rows $(1, 1, 2, 3, 5), (2, 4, 4), (3, 5, 5), (6, 7)$. Perhaps a better way to visualize this process is with the following notation:

$$\begin{array}{ccccccc} 1 & 1 & 2 & \underline{4} & 5 & \leftarrow & 3 \\ 2 & 4 & \underline{5} & & & \leftarrow & \\ 3 & 5 & \square & & & \leftarrow & \\ 6 & 7 & & & & & \end{array}$$

Here, the underlined entries of the tableau T are those elements which are displaced under the successive row insertions, and the box indicates the position at which row insertion stops because the new element is appended to a row of T. The positions of the underlined elements and the box form the *bumping path of t into T*.

The key lemma upon which the correctness of the Robinson-Schensted correspondence rests is the following.

Proposition 4.24. (a) *Let R be a row, and let $t_1 \leq t_2$ be positive integers. Perform the row insertions $(t_1'|\,R') \Longleftarrow (R|\,t_1)$ followed by $(t_2'|\,R'') \Longleftarrow (R'|\,t_2)$. Then t_2 is inserted into a position in R' strictly to the right of the position in R into which t_1 is inserted. Furthermore, either $t_2' = 0$ or $t_1' \leq t_2'$.*
(b) *Let T be a tableau with two rows R_1 and R_2. For $t_1 \in \mathbb{P}$, perform the row insertions $(t_2|\,R_1') \Longleftarrow (R_1|\,t_1)$ followed by $(t_3|\,R_2') \Longleftarrow (R_2|\,t_2)$. Either $t_2 = 0$ or t_2 is inserted into a position in R_2 weakly to the left of the position in R_1 into which t_1 is inserted.*

Proof. For part (a), since $t_1 \leq t_2$, it follows easily from the definition of row insertion that t_2 is inserted into R' strictly to the right of t_1. For the remaining claim we may assume that $t_1' \neq 0$ and $t_2' \neq 0$. Now t_1' and t_2' are the entries of R which formerly occupied the positions in R'' of t_1 and t_2, respectively. Since t_2 is strictly to the right of t_1, it follows that $t_1' \leq t_2'$.

For part (b), we may assume that $t_2 \neq 0$. Assume that t_1 is inserted into position b of R_1, displacing t_2. Then either position b of R_2 is empty, or t_2 is strictly less than the element in position b of R_2 (since the columns of T are strictly increasing). By definition of row insertion, it follows that t_2 is inserted into one of positions $1, 2, \ldots, b$ of R_2. $\square$

Applying Proposition 24 inductively, the proof of the following corollary is immediate.

Corollary 4.25. *Let T be a tableau, and let $t_1 \leq t_2$ be positive integers. Perform the tableau insertions $T' := (T \leftarrow t_1)$ followed by $T'' := (T' \leftarrow t_2)$.*

(a) *The bumping path of t_1 into T moves weakly to the left as one reads the successive rows of T. (Similarly for the bumping path of t_2 into T', of course.)*
(b) *The bumping path of t_2 into T' stays strictly to the right of the bumping path of t_1 into T.*

To illustrate Corollary 25, consider the following example:

$$\begin{array}{cccccccc}
1 & 1 & 2 & 3 & \underline{5} & 6 & \leftarrow & 4 \\
2 & 3 & \underline{6} & 7 & & & \leftarrow & \\
4 & 4 & \underline{7} & 8 & & & \leftarrow & \\
5 & 6 & \underline{8} & & & & \leftarrow & \\
\square & & & & & & \leftarrow &
\end{array}$$

followed by

$$\begin{array}{cccccccc}
1 & 1 & 2 & 3 & 4 & \underline{6} & \leftarrow & 5 \\
2 & 3 & 5 & \underline{7} & & & \leftarrow & \\
4 & 4 & 6 & \underline{8} & & & \leftarrow & \\
5 & 6 & 7 & \square & & & \leftarrow & \\
8 & & & & & & &
\end{array}$$

The resulting tableau is

$$\begin{array}{cccccc}
1 & 1 & 2 & 3 & 4 & 5 \\
2 & 3 & 5 & 6 & & \\
4 & 4 & 6 & 7 & & \\
5 & 6 & 7 & 8 & & \\
8 & & & & &
\end{array}$$

Proposition 4.26. *Let $t_1, t_2, \ldots, t_n$ be any sequence of positive integers, let $T_0 := ()$, and for $1 \le u \le n$ let $T_u := (T_{u-1} \leftarrow t_u)$. Then T_n is a tableau.*

Proof. By induction on n, it suffices to show that if T is a tableau and $t \in \mathbb{P}$, then $T' := (T \leftarrow t)$ is a tableau. By the definition of the row insertion algorithm, each row of T' is weakly increasing. Since the bumping path of t into T moves weakly to the left, the shape of T' is the Ferrers diagram of a partition obtained by adjoining one new box to the shape of T. Now, let's check the column condition for T'; so consider two consecutive entries $T'(a, b)$ and $T'(a + 1, b)$ in column b of T'. Since T is a tableau, we have $T'(a, b) < T'(a + 1, b)$ unless at least one of (a, b) or $(a + 1, b)$ is on the bumping path of t into T. Notice that (a, b) is not the new box of the shape of T'; if it were then $(a + 1, b)$ would be a box of T and the shape of T would not be the Ferrers diagram of a partition. If (a, b) is on the bumping path, then $T'(a, b)$ is strictly less than the element t' which is bumped out of row a of T', and t' is less than or equal to $T'(a + 1, b)$ since the bumping path is moving weakly to the left; in this case $T'(a, b) < T'(a+1, b)$, as required. If $(a + 1, b)$ is on the bumping path, then $T'(a + 1, b)$ is the value bumped out of row a of T'. Since the bumping path is moving weakly to the left, $T'(a, b)$ is less than or equal to the value t'' inserted into row a of T. Since t'' displaces $T'(a + 1, b)$ from row a of T, we have $t'' < T'(a + 1, b)$, completing the proof. $\square$

Iterating part (b) of Corollary 25, we obtain the following statement.

Proposition 4.27. *Let T_1 be a tableau, and let $t_1 \leq t_2 \leq \cdots \leq t_k$ be positive integers. Perform the tableau insertions $T_{i+1} := (T_i \leftarrow t_i)$ for each $1 \leq i \leq k$. Let λ be the shape of T_1 and let μ be the shape of T_{k+1}. Then $F_\mu \smallsetminus F_\lambda$ is a set of k boxes, with no more than one box per column.*

Because of this proposition, we can define a function from biwords β to pairs of tableau (T, U) of the same shape, as follows. Let $\beta = ((i_1, j_1), \ldots, (i_n, j_n))$ be a biword, and initially define $T_0 := ()$ and $U_0 := ()$ and $\lambda^{(0)} := ()$. For each $1 \leq u \leq n$, let $T_u := (T_{u-1} \leftarrow i_u)$ and let $\lambda^{(u)}$ be the shape of T_u. Construct U_u by adjoining to U_{u-1} a new box in the position $F_{\lambda^{(u)}} \smallsetminus F_{\lambda^{(u-1)}}$ which contains the value j_u. The *Robinson-Schensted correspondence* is the function $\beta \mapsto (T, U) := (T_n, U_n)$ defined by this procedure.

For example, consider the biword β pictured below, in which the first row is $i_1, \ldots, i_n$ and the second is $j_1, \ldots, j_n$.

$$\begin{array}{c|cccccccccccc} \mathbf{x} & 3 & 3 & 6 & 2 & 4 & 5 & 1 & 3 & 4 & 4 & 1 & 3 \\ \mathbf{y} & 1 & 1 & 1 & 2 & 2 & 2 & 3 & 3 & 3 & 3 & 4 & 4 \end{array}$$

Then (T_3, U_3), (T_6, U_6), (T_{10}, U_{10}), and $(T, U) = (T_{12}, U_{12})$ are as shown below.

$$\begin{array}{ccccc} 3 & 3 & 6 & & \end{array} \qquad \begin{array}{ccccc} 1 & 1 & 1 & & \end{array}$$

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

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

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

By Proposition 26, we know that T is a tableau; also, from Proposition 27 it follows that U is a tableau. It is clear from the construction that T and U have the same shape and that $(\mathbf{xy})^\beta = \mathbf{x}^T \mathbf{y}^U$ if $\beta \mapsto (T, U)$. All that remains is to prove that the Robinson-Schensted correspondence is a bijection from biwords to pairs of tableau of the same shape.

To see how we can reverse the above procedure $\beta \mapsto (T, U)$, consider the output pair (T, U) of tableau of the same shape F_λ. We must reconstruct the biword $\beta = ((i_1, j_1), \ldots, (i_n, j_n))$ corresponding to (T, U). First of all, the sequence $j_1, \ldots, j_n$ is just the sequence of entries of U sorted into weakly increasing order. Also from U, we can determine the sequence in which the boxes of F_λ were adjoined to the Ferrers diagram: they must appear in weakly increasing order of the entries of U, and among the boxes with equal entries of U they are adjoined from left to right, because of Corollary 25(b). For example, from the last tableau U pictured above, we see that the boxes of F_{5331} were adjoined in the order

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

while the pair (T, U) was being constructed. This allows us to reconstruct the sequence $i_1, \ldots, i_n$ of the biword β by "extracting" the entries of T according to the reverse order in which the boxes of F_λ were adjoined.

The tableau extraction algorithm is the functional inverse of the tableau insertion algorithm $(T \leftarrow t)$. It takes as input a tableau T and a box (a, b) of the shape of T which is a "lower-right corner" of the shape of T: that is, (a, b) is maximal in the product partial order ($(a', b') \leq (a, b)$ if and only if both $a' \leq a$ and $b' \leq b$) on the shape of T. The output is another tableau T' and a positive integer t; we denote this by $(T \stackrel{(a,b)}{\longrightarrow} t)$. To define tableau extraction, let the rows of T be $R_1, \ldots, R_k$, so that (a, b) is the last box of row a of the shape of T. Remove $t_a := T(a, b)$ from R_a along with the box (a, b) from the shape of T, and let R'_a denote the resulting row. In the next paragraph, we define a "row extraction" subroutine, denoted by $(t|\, R) \Longrightarrow (R'|\, t')$; this will be the functional inverse of row insertion. Then tableau extraction is defined by iterating row extraction as follows:

$$\begin{array}{rcl} (t_a|\, R_{a-1}) & \Longrightarrow & (R'_{a-1}|\, t_{a-1}) \\ (t_{a-1}|\, R_{a-2}) & \Longrightarrow & (R'_{a-2}|\, t_{a-2}) \\ & \vdots & \\ (t_3|\, R_2) & \Longrightarrow & (R'_2|\, t_2) \\ (t_2|\, R_1) & \Longrightarrow & (R'_1|\, t_1). \end{array}$$

The output is the tableau T' with rows $(R'_1, \ldots, R'_a, R_{a+1}, \ldots, R_k)$ and the positive integer $t := t_1$. (We must check that T' really is a tableau!)

The row extraction algorithm is defined as follows. Let $R: \ r_1 \leq \cdots \leq r_b$ be a row of positive integers, and let $t \geq 2$ be an integer. Let $1 \leq u \leq b$ be the greatest index such that $r_u < t$. Then $t' := r_u$ and $R': \ r_1 \leq \cdots \leq r_{u-1} < t \leq r_{u+1} \leq \cdots \leq r_b$. This defines row extraction $(t|\, R) \Longrightarrow (R'|\, t')$. By iteration, this defines tableau extraction $T \stackrel{(a,b)}{\to} t'$.

Proposition 4.28. *Let R be a weakly increasing sequence of positive integers.*
(a) *If $t \in \mathbb{P}$ and $(t'|\,R') \Longleftarrow (R|\,t)$ and $(t'|\,R') \Longrightarrow (R''|\,t'')$, then $R'' = R$ and $t'' = t$.*
(b) *If $t \geq 2$ is an integer and $(t|\,R) \Longrightarrow (R'|\,t')$ and $(t''|\,R'') \Longleftarrow (R'|\,t')$, then $R'' = R$ and $t'' = t$.*

Proof. Exercise. □

Proposition 4.29. *Let T be a tableau and let (a, b) be a lower-right corner of the shape of T. Perform the tableau extraction $T' := (T \xrightarrow{(a,b)} t)$. Then T' is a tableau.*

Proof. Exercise. □

From the above results it follows that the function $\beta \mapsto (T, U)$ defined above provides a weight-preserving bijecton between the set of all biwords β and the set of all pairs of tableau (T, U) of the same shape. This completes the verification of the correctness of the Robinson-Schensted correspondence, and proves Proposition 23.

One consequence of Proposition 23(a) is particularly noteworthy. For a partition $\lambda \vdash n$, a *standard Young tableau* of shape λ (abbreviated "SYT") is a tableau T of shape λ and content 1^n; that is, the entries of T are $\{1, 2, \ldots, n\}$. Let f_λ denote the number of SYTs of shape λ.

Exercise: By extracting the coefficient of $x_1 x_2 \cdots x_n y_1 y_2 \cdots y_n$ from both sides of Proposition 23(a), we see that

$$n! = \sum_{\lambda \vdash n} f_\lambda^2.$$

THE HOOK-LENGTH FORMULA

Let λ be a partition, and let F_λ be the corresponding Ferrers diagram. For any $(a, b) \in F_\lambda$, let the *hook of* (a, b) be the set of boxes $(a', b') \in F_\lambda$ such that either $(a' = a$ and $b' \geq b)$ or $a' \geq a$ and $b' = b)$; define $h(a, b)$ to be the cardinality of the hook of (a, b).

Proposition 4.30 (Frame, Robinson, and Thrall)**.** *Let λ be a partition. Then*

$$f_\lambda = \frac{n!}{\prod_{(a,b)\in F_\lambda} h(a, b)}.$$

Proof. See Section 5, Example 2 of Macdonald. □

SKEW SCHUR FUNCTIONS AND THE LITTLEWOOD-RICHARDSON RULE

Let λ and μ be partitions such that $F_\mu \subseteq F_\lambda$. A *skew tableau of shape* λ/μ is a function $T : F_\lambda \setminus F_\mu \to \mathbb{P}$ which is weakly increasing from left to right along rows, and strictly increasing from top to bottom along columns. The *skew Schur function of shape* λ/μ is $s_{\lambda/\mu} := \sum_T \mathbf{x}^T$, with the sum over all skew tableau of shape λ/μ.

Exercise: Sketch a proof that $s_{\lambda/\mu}$ is a symmetric function.

Exercise: Derive a formula for $s_{\lambda/\mu}$ as a polynomial in the complete symmetric functions.

Proposition 4.31. *Let λ and μ be partitions such that $F_\mu \subseteq F_\lambda$. The skew Schur function $s_{\lambda/\mu}$ is a nonnegative integer linear combination of ordinary Schur functions. That is,*

$$s_{\lambda/\mu} = \sum_{\nu} c^{\lambda}_{\mu\nu} s_\nu,$$

in which each $c^{\lambda}_{\mu\nu}$ is a nonnegative integer.

Proof. See Section 9 of Macdonald. □

In fact, to be more precise, the *Littlewood-Richardson coefficients* $c^{\lambda}_{\mu\nu}$ can be given a direct combinatorial interpretation, but we've done enough of symmetric functions and will move on to other things instead.