LTCC Enumerative Combinatorics

Notes 8 | Alex Fink | Fall 2015

8 Symmetric functions

The algebra of symmetric functions makes appearances in a wide range of problems, combinatorial and otherwise. Within combinatorics, a strength of theirs is problems involving integer partitions. We give a partial introduction here.

Note that these notes have material in common with the end of Section 4. The present treatment is less harried than that one. (The duplication is something I'll have to set right in the next instance of the course.)

8.1 The ring of symmetric functions and its bases

Let R be a divisible commutative ring: $\mathbb{Q}$ or $\mathbb{C}$ or similar are fine choices. The polynomial ring $S = R[x_1, \ldots, x_n]$ bears an action of the symmetric group $\mathfrak{S}_n$ which permutes the variables. To be explicit, for $\sigma \in \mathfrak{S}_n$, we define $\sigma \cdot x_i = x_{\sigma(i)}$; this defines the action completely once we impose the condition that $f \mapsto \sigma f$ be an R-algebra homomorphism for each σ.

The *ring of symmetric functions* in n variables,

$$\Lambda^n := R[x_1, \ldots, x_n]^{\mathfrak{S}_n},$$

is defined to be the subring of polynomials fixed by the action of every permutation in $\mathfrak{S}_n$. I will suppress mention of R in the notation.

It will be useful to have a concise notation for monomials in S. Given a vector $a = (a_1, \ldots, a_n) \in \mathbb{N}^n$, we write

$$x_a := x_1^{a_1} \cdots x_n^{a_n}.$$

We also allow the notation x^A where A is a subset of or multiset on $[n]$. These sets and multisets stand for vectors in the spirit of Section 2.4. So if $a \in \mathbb{N}^n$ and A is a multiset on $[n]$ containing i with multiplicity a_i for each $i \in [n]$, we again write $x^A := x^a$. Note that A is a set in the case that $a \in \{0,1\}^n$.

For any $\sigma \in \mathfrak{S}_n$ and any monomial $m \in R[x_1, \ldots, x_n]$, we have that $\sigma \cdot m$ is also a monomial. So a polynomial f is invariant under σ if and only if certain *equalities* hold between pairs of its coefficients. The conditions for f to be in Λ^n are therefore of the same form. To wit, if x^a and x^b are two monomials in S, their coefficients in each symmetric function agree if and only if their exponent vectors

a and b are permutations of one another. It follows that, for every nondecreasing exponent vector $\lambda \in \mathbb{N}^n$ with $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$, the polynomial

$$m_\lambda = \sum_{a \text{ is a permutation of } \lambda} x^a$$

is a symmetric function. We call these m_λ the *monomial symmetric functions*.

It is standard to speak of the indexing objects as (integer) *partitions*. We define

$$\mathrm{Par}^n := \{\lambda \in \mathbb{N}^n : \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n\}$$

to be the set of partitions with at most n parts: it is "at most" because we have allowed parts to equal zero, which we did not do in the definition of integer partition from Section 2.3. The above discussion establishes that

Proposition 8.1 *The set $\{m_\lambda : \lambda \in \mathrm{Par}^n\}$ is an R-module basis for Λ^n.*

It also follows that the sum of the degree k terms in a symmetric function is also a symmetric function. That is, the ring of symmetric functions is a graded ring, inheriting its grading from S where $\deg(x_i) = 1$ for all i. If T is a graded R-module (including an R-algebra, like S or Λ^n) and $k \in \mathbb{N}$, we denote by T_k the R-module of all elements of T homogeneous of degree k.

8.1.1 Integer partitions

We will need to introduce a few operations on integer partitions. First of all, as we have done above in the discussion of Par^n, we identify two nonincreasing sequences of naturals as being the same partition if they differ only by one of them having extra zeroes at the end. For example, this identification makes Par^m a subset of Par^n when $m < n$. This lets us define the set of all integer partitions as $\mathrm{Par} = \bigcup_{n \geq 0} \mathrm{Par}^n$.

We will sometimes abbreviate successive identical parts in a partition with a superscript: $(\ldots, a^k, \ldots)$ stands for $(\ldots, \overbrace{a, \ldots, a}^{k}, \ldots)$.

The *Ferrers diagram* or *Young diagram* of a partition λ is the subset

$$\{(i, j) \in (\mathbb{N}_+)^2 : i \in [\lambda_j]\}$$

of $(\mathbb{N}_+)^2$. This diagram is drawn as a set of boxes in the plane. There are multiple conventions as to how this is done, but as I'm teaching in English I'll use the *English notation* here. In the English notation, the pair (i, j) is drawn as a box at the point $(i, -j)$, so that λ_1 is the length of the *top* row of the diagram. For example, here is the Young diagram of the partition $(8, 4, 3, 1, 1)$.

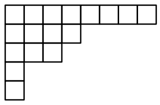

The *French notation* is the reflection of the English notation across a horizontal axis, i.e. with the pair (i,j) drawn at the point (i,j), so λ_1 measures the bottom row. You may also encounter a *Russian notation* where the whole contraption is tilted diagonally, with $(1,1)$ bottommost.

Via Young diagrams, we see that integer partitions are in bijection with finite downsets of the poset $\mathbb{N}\times\mathbb{N}$ (to which $\mathbb{N}_+\times\mathbb{N}_+$ is isomorphic). That is, a partition is a set of "boxes shoved into the corner of a room", with no gaps allowing any box to be pushed in horizontally or vertically closer. The bijection preserves size: if the Young diagram contains d elements then $\lambda_1+\lambda_2+\cdots=d$. In this case we say that λ is a *partition **of** d* and write $\lambda\vdash d$, or $|\lambda|=d$. We write Par_d for the set of partitions of d, and Par_d^n for the set of those with at most n parts[1].

The involutive automorphism of $\mathbb{N}\times\mathbb{N}$ which switches the two factors induces an involution on partitions, called *conjugation*. The conjugate of a partition λ is written λ'; explicitly,

$$\lambda_i' = \max\{j : \lambda_j \geq i\}.$$

For example, the conjugate of the partition $(8,4,3,1,1)$ depicted above is $(5,3,3,2,1,1,1,1)$. If $\lambda\vdash d$ then also $\lambda'\vdash d$.

The *dominance* or *majorisation* order on Par_d is the partial order defined by $\lambda\leq\mu$ if and only if

$$\lambda_1+\cdots+\lambda_i \leq \mu_1+\cdots+\mu_i$$

for each i.

8.1.2 First examples of symmetric functions

Historically, one of the first contexts in which symmetric functions were investigated was the study of roots of polynomials. The "general" polynomial

$$(y-x_1)\cdots(y-x_n)\in\Lambda^n[y]$$

is manifestly fixed under the $\mathfrak{S}_n$-action, since permutations $\sigma\in\mathfrak{S}_n$ act by permuting the factors. The coefficient of y^{n-k} in its expansion is the monomial symmetric function m_λ where $\lambda=(\underbrace{1,\ldots,1}_{k},\underbrace{0,\ldots,0}_{n-k})$. We give this symmetric function the special notation e_k. Explicitly,

$$e_k = \sum_{A\subseteq[n]\,:\,|A|=k} x^A.$$

[1]Stanley and those who follow him use $\mathrm{Par}(d)$. My notation is by analogy with the notation for components of a graded module, which has greater claim to being a standard.

The e_k are called the *elementary symmetric functions*.

Similarly given names are the *complete homogeneous symmetric functions*

$$h_k = \sum_{A \text{ a multiset on } [n] : |A|=k} x^A$$

and the *power sum symmetric functions*

$$p_k = \sum_{i \in [n]} x_i^k$$

which equal $m_{(k,0,\dots,0)}$ when $k > 0$. The above definitions dictate in particular that $e_k = h_k = 0$ when $k < 0$. (We won't need the p_k of negative index.)

These families have ordinary generating functions in $\Lambda^n[\![t]\!]$:

$$\begin{aligned}
\sum_k e_k t^k &= \prod_{i \in [n]} (1 + x_i t), \\
\sum_k h_k t^k &= \prod_{i \in [n]} \frac{1}{1 - x_i t}, \\
\sum_{k \geq 0} p_k t^k &= \sum_{i \in [n]} \frac{1}{1 - x_i t}.
\end{aligned}$$

The first two of these can be interpreted as "finely weighted" generating functions for subsets and multisets, where each element $i \in [n]$ gets its own weighting variable x_i, and in this way they generalise the generating functions of Section 2. One manifestation of our reciprocity between subsets and multisets is the fact that the first two of these generating functions become inverses once $-t$ is substituted for t in one of them:

$$\left(\sum_k e_k (-t)^k \right) \left(\sum_k h_k t^k \right) = \prod_{i \in [n]} \frac{1 - x_i t}{1 - x_i t} = 1.$$

Extracting coefficients of powers of t implies

$$h_n - e_1 h_{n-1} + e_2 h_{n-2} - \cdots + (-1)^n e_n = \begin{cases} 1 & n = 0 \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

If $\lambda = (\lambda_1, \dots, \lambda_\ell)$ is a partition (note that the upper index may exceed n) then we write $e_\lambda = \prod_{i=1}^{\ell} e_i$ and $h_\lambda = \prod_{i=1}^{\ell} h_i$ and $p_\lambda = \prod_{i=1}^{\ell} p_i$. Note that if $\lambda \vdash d$, then e_λ, h_λ, and p_λ, as well as m_λ, are homogeneous polynomials of degree d.

Theorem 8.2 *The set $\{e_\lambda : \lambda' \in \mathrm{Par}^n\}$ is an R-module basis for Λ^n.*

Proof It is enough to show that, for each $d \geq 0$, those e_λ with $\lambda' \in \mathrm{Par}^n_d$ constitute an R-module basis for Λ^n_d. We already know one basis for this module, consisting of the m_λ with $\lambda \in \mathrm{Par}^n_d$, by Proposition 8.1. Our strategy will be to analyse the matrix expressing each of these e_λ as a $\mathbb{Z}$-linear combination of the m_λ, and show it is triangular and hence invertible under a suitable order on Par^n_d, which will be majorisation (reading the indices of the e_λ in conjugate sense). It is worth stating the entries of this matrix as a lemma.

Lemma 8.3 *We have*

$$e_\lambda = \sum_\mu M_{\lambda\mu} m_\mu$$

where $M_{\lambda\mu}$ is the number of subsets S of $(\mathbb{N}_+)^2$ so that for each $i \in \mathbb{N}_+$ there are λ_i elements of S with first coordinate i, and for each $j \in \mathbb{N}_+$ there are μ_j elements of S with second coordinate j.

Proof The terms in the expansion of the product

$$e_\lambda = e_{\lambda_1} \cdots e_{\lambda_\ell} = \prod_{i\in[\ell]} \sum_{A\subseteq[n]\,:\,|A|=\lambda_i} x^A,$$

are indexed by ℓ-tuples of sets $A_i \subseteq [n]$ with $|A_i| = \lambda_i$. If such a term appears in the monomial symmetric function m_μ, then writing $x^{A_1} \cdots x^{A_\ell} = x^b$ for some $b \in \mathbb{N}^n$, the vector b is a permutation of μ. Within each copy of m_μ there will be just one term where b equals μ without permuting. In this case $S = \{(i, j) \in (\mathbb{N}_+)^2 : j \in A_i\}$ is one of the objects counted by $M_{\lambda\mu}$. This correspondence is bijective, and the lemma follows.

Proof of Theorem 8.2, resumed Let $\mu \in \mathrm{Par}^n_d$. It is clear that, to make λ as large as possible in the dominance order subject to $M_{\lambda\mu} \neq 0$, one should take the first coordinates of elements of S as small as possible. The smallest possibility of all for these first coordinates occurs when S is the Young diagram of μ, in which case $\lambda = \mu'$. Moreover this is the unique way to obtain $\lambda = \mu'$. So the matrix of $M_{\lambda\mu}$ is triangular with 1s on the diagonal, and the proof goes through.

Theorem 8.2 implies

Corollary 8.4 *The ring $\Lambda^n \cong R[e_1, \ldots, e_n]$ is a polynomial ring in n generators, of degrees* $\deg(e_i) = i$.

In particular, every symmetric (polynomial) function in the roots of a polynomial is a (polynomial) function of its coefficients.

This corollary was also, I'd argue, the first major result in classical *invariant theory*, which is concerned with the subalgebras S^G of invariants in the polynomial ring S under the action of a group G. Noether proved that S^G is finitely generated if G is finite. But it is not usually a polynomial ring again. The Chevalley-Shepard-Todd theorem states that, over $\mathbb{C}$, if the action of G on S preserves the grading, then S^G is a polynomial ring exactly when G acts on S_1 as a *complex reflection group*, i.e. G has a generating set each of whose elements fixes a codimension 1 subspace of S_1. If G needn't be finite then the situation is worse: Nagata's example of a non-finitely-generated subalgebra of a polynomial ring, which answered Hilbert's fourteenth problem in the negative, is the ring of invariants of a linear algebraic group.

8.1.3 Symmetric functions in infinitely many variables

Given naturals $m < n$, there is an inclusion of rings $\iota : \Lambda^m \hookrightarrow \Lambda^n$ which sends $e_k \in \Lambda^m$ to $e_k \in \Lambda^n$ for all $k \leq m$. Note that this is not simply the restriction of the usual inclusion $R[x_1,\ldots,x_m] \hookrightarrow R[x_1,\ldots,x_n]$; no nonconstant symmetric functions are in the image of this latter inclusion.

We can thus define *the* graded R-algebra of symmetric functions Λ to be the direct limit of the inclusions ι, that is the union of all the Λ^n identified under these inclusions. Then Λ is a polynomial ring in countably many generators, $\Lambda = R[e_1,e_2,\ldots]$, with $\deg e_i = i$. Informally, an element of Λ is a symmetric polynomial "in infinitely many variables".

The rings Λ^n also bear a family of projections, $\pi : \Lambda^n \twoheadrightarrow \Lambda^m$ for $n > m$, given by restriction of the usual projections on the polynomial rings, $\pi(x_i) = x_i$ for $i \leq m$ and $\pi(x_i) = 0$ for $i > m$. These satisfy the condition that

$$\Lambda^m \overset{\iota}{\hookrightarrow} \Lambda^n \overset{\pi}{\twoheadrightarrow} \Lambda^m$$

is the identity map on Λ^m, for all $m < n$. This means that given any R-algebra A and elements $a_1,a_2,\ldots \in A$, all but finitely many of which are zero, there is an evaluation map $\Lambda \to A$ which substitutes a_i for x_i.

Proposition 8.5 *Each of* $\{m_\lambda : \lambda \in \mathrm{Par}\}$*,* $\{e_\lambda : \lambda \in \mathrm{Par}\}$*,* $\{h_\lambda : \lambda \in \mathrm{Par}\}$*, and* $\{p_\lambda : \lambda \in \mathrm{Par}\}$ *is an R-module basis for* Λ.

For the m_λ and e_λ this follows from the previous sections. For the h_λ, equation (1) can be used recursively to express each e_i as a polynomial in the h_j, proving that the h_λ generate Λ as an R-module. Since the leading term in dominance order of e_i is $h_{(i)}$, that of e_λ is h_λ, so this linear transformation is upper-triangular in each graded component, showing independence of the $\{h_\lambda\}$. We leave the proof for the p_λ as an exercise.

Note also that, if the algebra morphism $\omega : \Lambda \to \Lambda$ is defined by $\omega(e_k) = h_k$, the symmetry of equation (1) implies that $\omega(h_k) = e_k$. That is, ω is an involution.

8.2 Schur polynomials

The nicest basis of all for Λ is that of the *Schur polynomials*. Given a partition $\lambda \in \mathrm{Par}^n$, define

$$s_\lambda = \frac{\begin{vmatrix} x_1^{\lambda_1+n-1} & x_2^{\lambda_1+n-1} & \cdots & x_n^{\lambda_1+n-1} \\ x_1^{\lambda_2+n-2} & x_2^{\lambda_2+n-2} & \cdots & x_n^{\lambda_2+n-2} \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{\lambda_n} & x_2^{\lambda_n} & \cdots & x_n^{\lambda_n} \end{vmatrix}}{\prod_{i<j\in[n]} (x_i - x_j)}.$$

We claim that $s_\lambda \in \Lambda^n$. Since a determinant changes sign on transposal of two columns, the numerator $f = \det(x_j^{\lambda_i+n-i})_{i,j\in[n]}$ of s_λ is an *alternating polynomial*: that is, for $\sigma \in \mathfrak{S}_n$, we have $\sigma f = (-1)^\sigma f$, where $(-1)^\sigma$ is 1 if σ is an even permutation and -1 if it is odd. (Said otherwise, there are two one-dimensional representations of $\mathfrak{S}_n$, the trivial and the sign representations; just as a symmetric polynomial generates a copy of the trivial representation, an alternating polynomial generates a copy of the sign representation). The denominator of s_λ is also alternating, since the transposition $(i\ i+1)$ negates one factor and permutes the others. So s_λ is at least a symmetric rational function.

Now, equating any two of the indeterminates x_i makes f vanish. So, viewed as a polynomial in x_n, each value $x_n = x_i$ for $i \in [n-1]$ is a root of f, and f is hence divisible by $\prod_{i\in[n-1]}(x_i - x_n)$. Induction on n implies that f is a polynomial.

Note that s_λ is homogeneous of degree d if $\lambda \vdash d$, for the numerator has degree $d + \binom{n}{2}$ and the denominator has degree $\binom{n}{2}$. For example, setting λ to the empty partition, $s_{()}$ must be a polynomial of degree zero, i.e. a scalar. The scalar in question is $s_{()} = 1$, as we see by comparing coefficients of $x_1^{n-1}x_2^{n-2}\cdots x_n^0$. Thus we have computed the determinant of the *Vandermonde matrix*:

$$\begin{vmatrix} x_1^{n-1} & \cdots & x_n^{n-1} \\ \vdots & \ddots & \vdots \\ x_1^0 & \cdots & x_n^0 \end{vmatrix} = \prod_{i<j\in[n]} (x_i - x_j).$$

Theorem 8.6 (Jacobi-Trudi identity) *We have*

$$s_\lambda = \det(h_{\lambda_i+j-i})_{i,j\in[n]}.$$

In particular, the $h_k = s_{(k)}$ are the Schur functions of one-row partitions. It also turns out that the $e_k = s_{(1^k)}$ are the Schur functions of one-column partitions: this can be shown by expanding the Jacobi-Trudi determinant along the k-th row, which by induction reduces to the identity (1).

Our proof is after Sagan.

Proof For $j \in [n]$, let $\iota_j : \Lambda^{n-1} \to S$ be the R-algebra map that sends the variables $x_1, \ldots, x_{n-1}$ respectively to $x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n$, omitting x_j. This omission transforms the generating function to

$$\sum_k \iota_j(e_k) t^k = \prod_{i \in [n] \setminus \{j\}} (1 + x_i t),$$

so

$$\left(\sum_k \iota_j(e_k)(-t)^k \right) \left(\sum_k h_k t^k \right) = \frac{1}{1 - x_j t}.$$

implying

$$h_n - \iota_j(e_1) h_{n-1} + \iota_j(e_2) h_{n-2} - \cdots + (-1)^n \iota_j(e_n) = x_j^n.$$

Define $n \times n$ matrices over S by $E = ((-1)^{n-1} \iota_j(e_{n-i}))_{i,j \in [n]}$ and, for any vector $a \in \mathbb{N}^n$, by $H(a) = (h_{a_i + n - j})_{i,j \in [n]}$. Then in the matrix product $H(a) \cdot E$ every entry is an alternating sum like the one displayed above, so

$$H(a) \cdot E = (x_j^{a_i})_{i,j \in [n]}.$$

We conclude

$$\begin{aligned} s_\lambda &= \frac{\det(H(\lambda_1 + n - 1, \ldots, \lambda_n + 0) \cdot E)}{\det(H(n-1, \ldots, 0) \cdot E)} \\ &= \frac{\det H(\lambda_1 + n - 1, \ldots, \lambda_n + 0)}{\det H(n-1, \ldots, 0)} \\ &= \det H(\lambda_1 + n - 1, \ldots, \lambda_n + 0), \end{aligned}$$

since $H(n-1, \ldots, 0)$ is triangular with ones along the diagonal, and this is the theorem.

8.2.1 Lindström-Gessel-Viennot and Young tableaux

The next lemma provides an interpretation of determinantal formulae in terms of families of paths. It was first proved by Lindström. The contribution of Gessel and Viennot was to notice its great combinatorial utility. It is, for instance, the underlying reason for the positivity of the classes constructed by Lascoux in his *Classes de Chern d'un produit tensoriel.*

We state a weighted version. Let G be a directed graph with no cycles (it is possible to relax this assumption, but we won't) with a weight $w(e) \in R$ on each edge e. The weight $w(S)$ of a subset S of the edges of G is the product $\prod_{e \in S} w(e)$. Say that a *routing* from an ordered list of *source* vertices $s = (s_1, \ldots, s_n)$ to an ordered list of *sink* vertices $t = (t_1, \ldots, t_n)$ consists of a finite directed path P_i from s_i to t_i for each i, visiting pairwise disjoint sets of vertices. Let $r(s,t)$ denote the sum of the weights of all routings from s to t. If $n = 1$ then a routing is just a directed path, so $r(s_i, t_j)$ is the total weight of paths from i to j. As usual with weighted enumeration, setting $w(e) = 1$ for all edges will turn $r(s,t)$ into a simple count of routings.

Lemma 8.7 (Lindström-Gessel-Viennot) *Let $s = (s_1, \ldots, s_n)$ and $t = (t_1, \ldots, t_n)$ be lists of vertices of G. Then*

$$\det(r(s_i,t_j))_{i,j\in[n]} = \sum_{\sigma\in\mathfrak{S}_n} (-1)^\sigma r(s,\sigma t).$$

In many applications G has a planar embedding in a di-gon with the vertices s lined up in order along the left side, and the vertices t lined up in order along the right side. Then topology rules out the existence of routings from s to non-identity permutation of t, and the lemma says $\det(r(s_i,t_j))_{i,j\in[n]} = r(s,t)$.

Proof Our proof is another application of the method of a sign-reversing involution, seen earlier in the proof of Proposition 3.5.

By Proposition 5.6 applied to the relation "there exists a directed path from v to w", there exists a total order on the vertices of G such that $v < w$ whenever (v,w) is an edge of G. Fix such an order.

Consider the set $P(s,t)$ of all n-tuples of paths from the vertices s in the given order to the vertices t in some order, weighted by the sign of the associated permutation (times the product of the edge weights). Then the left side of the equation of the lemma is the total weight of all elements of $P(s,t)$. The right side is the total weight of those tuples of paths in $P(s,t)$ which visit pairwise-disjoint vertex sets. The remaining tuples of paths cancel in pairs. Suppose $(P_1, \ldots, P_n)$ is a tuple of paths that is not vertex-disjoint, and v is the greatest vertex shared by at least two of its paths, say P_i, P_j, and some others of index greater than j. Construct two new paths made by cutting and pasting at j: let P_i' be the segment of P_i up to v fused with the segment of P_j from v on, and P_j' be similar with the roles of i and j reversed. Then replacing P_i and P_j by P_i' and P_j' gives a tuple of paths in $P(s,t)$ with equal weight but opposite sign.

Let λ be a partition. A *semi-standard Young tableau T* (*SSYT* for short) of shape λ is a set function from the Young diagram of λ to $[n]$, or to $\mathbb{N}_+$ if we work with infinitely many variables, which is weakly increasing along each row and strictly increasing along each column, i.e. with

$$T(i,j) \leq T(i+1,j)$$

and

$$T(i,j) < T(i,j+1)$$

for each i,j such that both boxes are in the Young diagram. Young tableaux are drawn by writing the integer $T(i,j)$ inside the box representing the element (i,j) of the Young diagram.

For example, SSYTs of shape (k) are in bijection with multisets on $[n]$ of size k, and SSYTs of shape (1^k) are in bijection with subsets of $[n]$ of size k, in both cases by listing the elements in nondecreasing order.

To each tableau T is associated a monomial x^T, equal to x^A where A is the multiset of values of T. This A, or its corresponding vector of naturals, is called the *content* of T.

Theorem 8.8 *Let λ be a partition. In Λ,*

$$s_\lambda = \sum_T x^T$$

where the sum ranges over SSYTs T of shape λ with codomain $\mathbb{N}_+$.

Proof It is enough to prove this in Λ^n for large enough n, where we work only with tableaux with codomain $[n]$.

Let G be the directed graph with vertices $\mathbb{Z} \times [n]$ and edges $(i,j) \to (i+1,j)$ of weight x_j ("horizontal" edges) and $(i,j) \to (i,j+1)$ of weight 1 ("vertical" edges). Choose the sources to be $s_i = (n-i,1)$ and the sinks to be $t_i = (\lambda_i + n - i, n)$, for $i \in [n]$. The total weight of the paths from s_j to t_i is $h_{\lambda_i + j - i}$, on account of the bijection between these paths and the multisets on $[n]$ of size $\lambda_i + j - i$, sending a path to the multiset of indices of the variables encountered as weights. There is exactly one path for each multiset, the one where the variables are encountered in nondecreasing order. So, by the Jacobi-Trudi identity and the planar digon case of the Lindström-Gessel-Viennot lemma, the total weight of routings from $(s_1, \ldots, s_n)$ to $(t_1, \ldots, t_n)$ equals s_λ.

These routings are in bijection with SSYTs of shape λ, by extending the idea of the last bijection. Given a routing, construct a tableau of shape λ by sending the Young diagram box (i,j) to the index of the ith variable encountered on the path from s_j to t_j. Then each row is nondecreasing as above, while each column is strictly increasing because vertex-disjointness of the paths demands that if the jth path contains a vertex (i,k) as the tail of a horizontal edge and the $j+1$st a vertex (i,k') as the head of one, then $k' > k$.

Proposition 8.9 *The set $\{s_\lambda : \lambda \in \mathrm{Par}^n\}$ is an R-module basis for Λ^n. Hence, $\{s_\lambda : \lambda \in \mathrm{Par}\}$ is a basis for Λ.*

Proof Again, we argue in one homogeneous component at a time, examining the matrix expressing the s_λ in terms of the m_λ implicit in Theorem 8.8. Let us consider only the tableaux T of content μ for some partition μ (i.e. where permuting μ is not needed). If there exists a semistandard T Young tableau of shape λ and content μ, then λ dominates μ, because clearly T can only map boxes in the first k rows of the tableau to integers $\leq k$, and this is the dominance inequality. Moreover, there is exactly one SSYT of shape λ and content λ, namely the one with every box (i,j) labelled j. So our matrix is triangular and invertible, and the proposition follows.

Exercise Consider the graph with vertices $\mathbb{Z}^2$ and edges $(i,j) \to (i+1,j)$ and $(i,j) \to (i,j+1)$, where the latter still has weight 1 but the former has weight x_{i+j} (never mind that the index is sometimes nonpositive). Show that the elementary symmetric functions e_k can be realised as $r(v,w)$ for suitably chosen single vertices v and w. Use the Lindström-Gessel-Viennot lemma to prove that

$$s_{\lambda'} = \det(e_{\lambda_i+j-i})_{i,j\in[n]}.$$

Because ω exchanges the matrix in the exercise with the one in the Jacobi-Trudi identity, it is a corollary of the exercise that $\omega(s_\lambda) = s_{\lambda'}$.

8.2.2 Plane partitions

As an enumerative application, we find a generating function for plane partitions. A *plane partition* is a finite downset of the poset $\mathbb{N} \times \mathbb{N} \times \mathbb{N}$. That is, it is a configuration of boxes stacked in the corner of a three-dimensional room. The name reflects that plane partitions are in bijection with infinite two-dimensional arrays of natural numbers, only finitely many nonzero, and nonincreasing in both directions, just as finite downsets of $\mathbb{N} \times \mathbb{N}$ are in bijection with integer partitions.

We will start by counting the downsets of a product of three finite total orders $\mathbf{a} \times \mathbf{b} \times \mathbf{c}$, and then let a and b and c tend to infinity. We proceed bijectively. Given such a downset, we pass to the associated two-dimensional arrays, whose (i,j)th entry is the number of k such that (i,j,k) is in the downset; this array can be taken as a map from the Young diagram of the partition (a^b) to $\{0,\dots,c\}$ which is nonincreasing as a and b increase. By subtracting the values from $c+1$, we get a map to $[c+1]$ which is nondecreasing as a or b increase. Finally, by adding $b-j$ to the entries of the j-th row, we get a SSYT for (a^b) with values in $[c+b]$.

If our original downset had cardinality n, then the corresponding nondecreasing array has weight x^n if we give each entry k the weight q^{c+1-k}. If we use a similar scheme for the SSYT, giving an entry k the weight q^{c+b-k}, the total weight is $q^{n+a\binom{b}{2}}$. Therefore, the generating function for our downsets by size is an evaluation of a Schur polynomial in $c+b$ variables, up to the extra factor of q, namely

$$d_{abc}(q) := q^{-a\binom{b}{2}}\, s_{(a^b)}(q^{c+b-1}, q^{c+b-2}, \ldots, 1).$$

The definition of this Schur function gives

$$d_{abc}(q) = q^{-a\binom{b}{2}} \frac{\det\big((q^{c+b-j})^{(a^b)_i + c+b-i} \big)_{i,j\in[c+b]}}{\prod_{0\le i<j<c+b} (q^j - q^i)}$$

in which the determinant at numerator is a Vandermonde matrix read transpose-wise, in the sequence of powers $q^{(a^b)_i+c+b-i}$, namely

$$q^{c+b+a-1}, q^{c+b+a-2}, \ldots, q^{c+a},\ q^{c-1}, q^{c-2}, \ldots, q^0.$$

So the numerator is the product $\prod q^j - q^i$ over all pairs of exponents $0 \le i < j < c+b+a$, excluding those where either i or j is among $c, c+1, \ldots, c+a-1$.

The factors of the numerator with $i, j < c$ cancel the corresponding factors of the denominator, while when $i, j \ge c+a$ the factor $q^j - q^i$ of the numerator is q^a times the factor $q^{i-a} - q^{j-a}$ of the denominator, so these $\binom{b}{2}$ factors can also be cancelled together with the $q^{-a\binom{b}{2}}$ at the front. What remains is

$$\begin{aligned} d_{abc}(q) &= \frac{\prod_{i=0}^{c-1} \prod_{j=c+a}^{c+b+a-1} (q^j - q^i)}{\prod_{i=0}^{c-1} \prod_{j=c}^{c+b-1} (q^j - q^i)} \\ &= \prod_{i=0}^{c-1} \prod_{j=c}^{c+b-1} \frac{1 - q^{j+a-i}}{1 - q^{j-i}} \end{aligned}$$

We reindex, replacing i by $c-1-i$ and j by $c+j$:

$$d_{abc}(q) = \prod_{i=0}^{c-1} \prod_{j=0}^{b-1} \frac{1 - q^{i+j+a+1}}{1 - q^{i+j+1}}. \tag{2}$$

This takes its most symmetric form when we un-telescope the inside.

Proposition 8.10 (MacMahon) *The generating function for downsets of* $\mathbf{a} \times \mathbf{b} \times \mathbf{c}$ *is*

$$d_{abc}(q) = \prod_{i=0}^{c-1} \prod_{j=0}^{b-1} \prod_{k=0}^{a-1} \frac{1 - q^{i+j+k+2}}{1 - q^{i+j+k+1}}.$$

Check that this agrees with our earlier methods for counting integer partitions when $c = 1$.

We can see that letting a tend to infinity simply kills the numerator of (2), while letting b and c tend to infinity makes the two products infinite. The number of times $i + j = n$ appears in the corresponding double product is $n + 1$. We conclude that

Proposition 8.11 *The generating function for plane partitions by size is*

$$\prod_{n\geq 1} \frac{1}{(1-q^n)^n}.$$

8.3 Looking beyond

I have had to cut these notes off much shorter than even many of the other sections. There are a whole wealth of combinatorial topics that belong here, not least the Robinson-Schensted-Knuth bijection between total orders on $[d]$ and pairs of identically-shaped d-box *standard* Young tableaux (that is, semi-standard Young tableau taking each value in $[d]$ once, or linear extensions of the Young diagram).

Instead, I'd like to point to some of the myriad appearances of the ring of symmetric functions in other areas. In many of these, the ring turns up furnished with its basis of Schur functions arising naturally, which is often detected by the appearance of the structure coefficients for their product. These are the *Littlewood-Richardson coefficients* $c^{\nu}_{\lambda\mu}$, defined by

$$s_\lambda s_\mu = \sum_\nu c^\nu_{\lambda\mu} s_\nu.$$

- The characteristic 0 irreducible representations ρ_λ of the symmetric group $\mathfrak{S}_d$ are labelled by partitions $\lambda \vdash d$. They multiply as Schur functions under the *induction product*: inducing $\rho_\lambda \otimes \rho_\mu$ from $\mathfrak{S}_d \times \mathfrak{S}_e$ to $\mathfrak{S}_{d+e}$ produces $\sum_\nu c^\nu_{\lambda\mu}\rho_\nu$.
- The characteristic 0 irreducible representations S_λ of the special linear group SL_n are labelled by partitions $\lambda \in \mathrm{Par}^{n-1}$. They multiply as Schur functions under the tensor product: $S_\lambda \otimes S_\mu = \sum_\nu c^\nu_{\lambda\mu} S_\nu$.

  The story can of course also be told on the Lie algebra level, and slight variations hold for related groups like general linear groups and special unitary groups. This item is related to the previous one by Schur-Weyl duality.
- The Grassmannian $Gr(n, \mathbb{C}^m)$ has an affine paving into *Schubert cells*, whose closures form a $\mathbb{Z}$-basis for its cohomology ring. They are indexed by partitions in Par^n whose conjugates lie in Par^{m-n}, that is, partitions whose Young

diagrams are subsets of $[m-n]\times[n]$. They multiply as Schur functions under the cup product in the cohomology ring.

- Let p be a prime. Up to isomorphism, finite abelian p-groups G_λ are indexed by partitions λ. The number of short exact sequences

  $$0\to G_\lambda\to G_\nu\to G_\mu\to 0,$$

  is a polynomial in p, which at $p=1$ specialises to $c^\nu_{\lambda\mu}$. In particular, such a sequence exists, independently of p, iff $c^\nu_{\lambda\mu}>0$.

  This extends to modules over any discrete valuation ring. It was one of the first-studied cases of the *Hall algebra*.

- An $n\times n$ Hermitian matrix has real eigenvalues; let us say that its spectrum is the list of these with multiplicity in nonincreasing order. For partitions $\lambda,\mu,\nu\in\mathrm{Par}^n$, there exists Hermitian matrices A,B, and $C=A+B$ with respective spectra λ, μ, ν if and only if $c^\nu_{\lambda\mu}>0$. In fact the volume of the space of such triples (A,B,C) is, up to a constant, the leading coefficient of $c^{m\nu}_{m\lambda,m\mu}$ as a polynomial in the scale factor m. As for spectra that need not be integers, the precise inequalities cutting out the cone of permissible triples of spectra is also described in terms of Littlewood-Richardson coefficients. (Describing these triples of spectra was known as the *Horn problem*.)

- The set of divisibility relations between the invariant factors of three square matrices over a commutative PID involve Littlewood-Richardson coefficients in a strikingly similar way to the above (see Thompson, *Divisibility relations satisfied by the invariant factors of a matrix product*, 1989).

- Schur functions make an appearance in the hierarchy of differential equations beginning with the KP equation, as labels of differential operators applied to a certain auxiliary function called the τ-function. I don't know whether the Littlewood-Richardson coefficient appear there in full array, but the Plücker relations among the Schur functions do, as do the inner product on the space of symmetric functions under which they are orthonormal. (See for instance work of Yuji Kodama.)

- Perhaps unsurprisingly there are appearances in physics: Wigner, *On the consequences of the symmetry of the nuclear Hamiltonian on the spectroscopy of the nuclei*, exploited the Littlewood-Richardson coefficients as early as 1937.

Rich theories can be drawn inter-relating nearly any pair of these appearances. A philosophy of Andrei Zelevinsky is that the more fundamental reason (as it were) for these relationships is simply that structural features of each of these problems demand that the rings be characteristic 0 Hopf algebras, indecomposible as tensor products, satisfying conditions of self-adjointness and positivity of their structure constants. The algebra Λ is (essentially) the unique such algebra, up to isomorphism.

Just as manifold as the settings the Littlewood-Richardson coefficients appear are the *Littlewood-Richardson rules* devised to compute them. Here are a few of my favourites.

- The classical Littlewood-Richardson rule says that $c^{\nu}_{\lambda\mu}$ counts the semistandard tableaux of a *skew partition* shape, that is the set difference of the Young tableaux of two partitions, with an extra condition on the sequence obtained by reading out the numbers in the cells in a certain order.

- The *puzzle rule* of Knutson and Tao says that $c^{\nu}_{\lambda\mu}$ counts the tilings of a large equilateral triangle with small equilateral triangles with their edges labelled by 0s and 1s in a restricted way, with the partitions λ, μ, and ν encoded by the edge labellings around the outside of the triangle.

  There is a mosaic rule of Kevin Purbhoo, introducing some further kinds of tiles, which connects these rules to the previous.

- Knutson and Tao also obtain the LR coefficients as counts of *hives*, these being plane graphs with all vertices trivalent with angles of $2\pi/3$, and allowing infinite rays in three of the six legal directions; λ, μ, and ν are encoded by the positions of the rays in each of these three directions.