

Lecture Notes for C&O 430/630
Fall 2010

Kevin Purbhoo¹

October 4, 2010

¹based in part on course materials written by I.P. Goulden and D.G. Wagner

Chapter 1

Formal power series

In this chapter, we develop the theory of formal power series and formal Laurent series, including Lagrange’s implicit function theorem and Hensel’s lemma. The main examples we’ll be considering in this course consider are generating functions. Before we get into the foundations, we’ll look at some examples of ordinary generating functions that satisfy certain functional equations, and the types of algebraic manipulations that can be used to solve them.

1.1 A first example: Catalan numbers

Let \mathcal{B} be the set of *binary trees*. These are rooted, ordered trees in which every vertex has up-degree 0 or 2. By rooted we mean that one vertex is specially marked (and generally drawn at the bottom). The up-degree of a vertex v is the number of up-edges incident with v , i.e. edges that are not on the path from v to the root vertex. Ordered means that we distinguish between the left and right branches above each vertex. The smallest binary tree consists of a single vertex, which is the root; we denote it by δ . Figure 1.1.1 shows all binary trees with at most five vertices:




Figure 1.1.1: Binary trees with at most 5 vertices.

We define a *weight function*, $\text{wt} : \mathcal{B} \rightarrow \mathbb{Z}$, as follows. For a binary tree $b \in \mathcal{B}$, $\text{wt}(b)$ is the number of vertices in b with updegree 2. Now form the generating function

$$B(x) = \sum_{b \in \mathcal{B}} x^{\text{wt}(b)}.$$

From Figure 1.1.1 we see that $B(x) = 1 + x + 2x^2 + \dots$. Collecting like terms, we can write $B(x) = \sum_{n \geq 0} b_n x^n$, where b_n is the number of binary trees of weight n . We say that b_n is the coefficient of x^n in $B(x)$ and write this as $b_n = [x^n]B(x)$.

$B(x)$ is an example of a formal power series. We'll say more about the foundations of formal power series starting in Section 1.3. For now, we'll treat them as (possibly mysterious) objects for which the usual rules of algebra hold: we can rearrange terms, collect like terms, use the distributive law to multiply, etc. In general, we allow ourselves to perform any operation we'd perform on polynomials or rational functions, provided any coefficient of the result is computable by a finite process. We do not allow ourselves to perform operations for which computing any coefficient involves an infinite sum or a limit in the traditional sense. For example $B(x)^2 = 1 + 2x + 5x^2 + \dots$ is legitimate, since $[x^n]B(x)^2$ depends only on $[x^k]B(x)$ for $k \leq n$. However, $B(B(x))$ is not a valid substitution, since the expression we obtain for $[x^0]B(B(x))$ is an infinite sum.

Almost every binary tree decomposes as the root vertex, plus an ordered pair of binary trees: namely, the left branch above the root and the right branch above the root. The only exception is the tree δ , which does not have any branches. Thus we have a recursive decomposition of \mathcal{B} , the set bijection

$$B \leftrightarrow \{\delta\} \sqcup \{\delta\} \times \mathcal{B} \times \mathcal{B}.$$

From this bijection, we obtain

$$\begin{aligned} B(x) &= \sum_{b \in \mathcal{B}} x^{\text{wt}(b)} \\ &= x^{\text{wt}(\delta)} + \sum_{b_1 \in \mathcal{B}} \sum_{b_2 \in \mathcal{B}} x^{1+\text{wt}(b_1)+\text{wt}(b_2)} \\ &= 1 + x^1 \left(\sum_{b_1 \in \mathcal{B}} x^{\text{wt}(b_1)} \right) \left(\sum_{b_2 \in \mathcal{B}} x^{\text{wt}(b_2)} \right) \\ &= 1 + xB(x)^2, \end{aligned}$$

a quadratic equation for $B(x)$.

One approach to solving this equation is to use the quadratic formula, which gives

$$B(x) = \frac{1 \pm \sqrt{1 - 4x}}{2x}.$$

From the binomial theorem

$$\sqrt{1 - 4x} = (1 - 4x)^{1/2} = 1 + \sum_{k \geq 1} \binom{\frac{1}{2}}{k} (-4)^k x^k,$$

where,

$$\begin{aligned} \binom{\frac{1}{2}}{k} (-4)^k &= \frac{\frac{1}{2} \cdot \frac{-1}{2} \cdot \frac{-3}{2} \cdots \frac{-2k+3}{2}}{k!} (-1)^k 2^k 2^k \\ &= -\frac{1 \cdot 3 \cdots (2k-3)}{k!} 2^k \\ &= -2 \frac{1 \cdot 3 \cdots (2k-3)}{k!} \frac{2 \cdot 4 \cdots (2k-2)}{(k-1)!} \\ &= -\frac{2}{k} \binom{2k-2}{k-1} \end{aligned}$$

and so substituting back in, we obtain

$$B(x) = \frac{1}{2x} \pm \left(\frac{1}{2x} - \frac{1}{x} \sum_{k \geq 1} \frac{1}{k} \binom{2k-2}{k-1} x^k \right).$$

However, the solution with a “+” sign has a leading term of $\frac{1}{x}$, which $B(x)$ does not have. Therefore the “-” solution is the correct one. Simplifying, we have

$$\begin{aligned} B(x) &= \sum_{k \geq 1} \frac{1}{k} \binom{2k-2}{k-1} x^{k-1} \\ &= \sum_{n \geq 0} \frac{1}{n+1} \binom{2n}{n} x^n, \end{aligned}$$

and so $b_n = \frac{1}{n+1} \binom{2n}{n}$ for $n \geq 0$. This is a classical sequence of combinatorial numbers called the *Catalan numbers*.

The Catalan numbers arise in a number of different combinatorial settings. Here are some examples of other combinatorial objects that are counted by Catalan numbers.

- Legitimate bracketings with n pairs of brackets. To be legitimate, the k^{th} left bracket must appear before the k^{th} right bracket, e.g. $()((()))$ is valid, but $((()))()$ is not.
- Weakly increasing sequences of positive integers $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ with $\lambda_i \leq i$.
- Rooted ordered trees with $n + 1$ vertices.

Exercise 1.1.1. For each of the combinatorial objects listed above, find a bijection with the set of binary trees of weight n .

A second approach to solving the quadratic equation

$$B(x) = 1 + xB(x)^2 \tag{1.1.1}$$

is to use Lagrange’s Implicit Function Theorem (LIFT). We’ll state it slightly imprecisely for now.

Theorem 1.1.2 (LIFT, slightly imprecise version). Let $\phi(\lambda)$, $f(\lambda)$ be formal power series, such that $[\lambda^0]\phi(\lambda)$ is invertible. Suppose $A(x)$ is a formal power series satisfying the function equation

$$A(x) = x\phi(A(x)).$$

Then for $n \geq 1$,

$$[x^n]f(A(x)) = \frac{1}{n}[\lambda^{n-1}]f'(\lambda)\phi(\lambda)^n.$$

In particular, if we take $f(\lambda) = \lambda$, we obtain the coefficients of $A(x)$:

$$[x^n]A(x) = \frac{1}{n}[\lambda^{n-1}]\phi(\lambda)^n.$$

LIFT does not work for $n = 0$; instead, it is not hard to see that $[x^0]f(A(x)) = [\lambda^0]f(\lambda)$, and in particular $[x^0]A(x) = 0$.

The quadratic equation (1.1.1) is not quite in the right form to apply LIFT. To get in the right form, we make the substitution $B(x) = A(x) + 1$, to obtain $A(x) = x(A(x) + 1)^2$. We now have the right form, with $\phi(\lambda) = (1 + \lambda)^2$. By LIFT, for $n \geq 1$ we have

$$\begin{aligned} b_n &= [x^n]B(x) = [x^n]A(x) \\ &= \frac{1}{n}[\lambda^{n-1}]1 \cdot (1 + \lambda)^{2n} \\ &= \frac{1}{n} \binom{2n}{n-1} \\ &= \frac{1}{n+1} \binom{2n}{n}, \end{aligned}$$

which is the same answer given by the quadratic formula.

Lagrange's theorem has some obvious advantages over the quadratic formula approach. For one thing, the second solution was significantly shorter. For another, it gives us information about related problems, without much more work. For example, the coefficients of $A(x)^k$ are given by

$$[x^n]A^k = \frac{1}{n}[\lambda^{n-1}]k\lambda^{k-1} \cdot (1 + \lambda)^{2n} = \frac{k}{n}[\lambda^{n-k}](1 + \lambda)^{2n} = \frac{k}{n} \binom{2n}{n-k}.$$

As we'll see, this can be useful in combinatorial contexts, when the generating function we're after is given as a function of the series for which we have a recursion. Of course the greatest advantage over the quadratic formula is that it applies to equations that are not quadratic. For example, if instead of considering binary trees, we consider trees in which each vertex has up-degree either 0 or m , a similar analysis gives the functional equation $C = x(1 + C)^m$ from which we obtain

$$[x^n]C^k = \frac{1}{n}[\lambda^{n-1}]k\lambda^{k-1} \cdot (1 + \lambda)^{mn} = \frac{k}{n}[\lambda^{n-k}](1 + \lambda)^{mn} = \frac{k}{n} \binom{mn}{n-k}.$$

1.2 A second example: rooted triangulations

A common usage of LIFT, which we'll need in this section, is to solve the following type of functional equation. Let $f(\lambda)$ and $g(\lambda)$ be given formal power series, and suppose that $H(x)$ is a formal power series satisfying

$$H(g(\lambda)) = f(\lambda).$$

We wish to solve for $H(x)$. To do this, let $\phi(\lambda) = \lambda/g(\lambda)$, and suppose $A(x)$ is a formal power series such that $A(x) = x\phi(A(x))$. Then

$$[x^n]H(x) = [x^n]H\left(\frac{A(x)}{\phi(A(x))}\right) = [x^n]H(g(A(x))) = [x^n]f(A(x)),$$

and using $[x^n]f(A(x)) = \frac{1}{n}[\lambda^{n-1}]f'(\lambda)\phi(\lambda)^n$, we find

$$[x^n]H(x) = \frac{1}{n}[\lambda^{n-1}]f'(\lambda)\frac{\lambda^n}{g(\lambda)^n} = [\lambda^{-1}]\frac{f'(\lambda)}{ng(\lambda)^n}.$$

Our next example involves an ordinary multivariate generating function, which is the type generating function we consider when we have more than one weight function on a set. The general set-up is as follows. Let \mathcal{S} be a set of combinatorial objects, and let $\text{wt}_1, \dots, \text{wt}_k : \mathcal{S} \rightarrow \mathbb{Z}$ be weight functions on \mathcal{S} . The generating function for \mathcal{S} with respect to these weight functions is

$$\Phi(x_1, \dots, x_k) = \sum_{\sigma \in \mathcal{S}} x_1^{\text{wt}_1(\sigma)} \dots x_k^{\text{wt}_k(\sigma)}.$$

If we collect like terms, we can write

$$\Phi(x_1, \dots, x_k) = \sum_{m_1, \dots, m_k} \phi_{m_1, \dots, m_k} x_1^{m_1} \dots x_k^{m_k},$$

where ϕ_{m_1, \dots, m_k} is the number of objects $\sigma \in \mathcal{S}$ with $(\text{wt}_1(\sigma), \dots, \text{wt}_k(\sigma)) = (m_1, \dots, m_k)$. Alternatively, we can think of $\text{wt} : \mathcal{S} \rightarrow \mathbb{Z}^k$, $\text{wt}(\sigma) = (\text{wt}_1(\sigma), \dots, \text{wt}_k(\sigma))$ as a single vector-valued weight function. From this point of view, it is not hard to see that the elementary counting lemmas (e.g. Sum Lemma, Product Lemma), are still valid.

Before proceeding, we pause for a quick aside on coefficient notation in the multivariate setting. If

$$f(x, y) = x^2 + 3x + 2xy^2 = x^2 + (3 + 2y^2)x,$$

it should be clear that $[xy^2]f(x, y) = 2$. However, what is meant by $[x]f(x, y)$ seems to be ambiguous. Both $[x]f(x, y) = 3$ and $[x]f(x, y) = 3 + 2y^2$ are valid interpretations. To resolve this, we will decide that this notation means the latter, and we'll write $[xy^0]f(x, y)$ when we mean the former. Be warned, however, that not every author may take this interpretation.

The goal of our second example is to count *rooted triangulations* in the plane. To begin, let \mathcal{T} be the set of rooted near-triangulations. These are connected graphs embedded in the plane, such that every interior face has degree 3. Multiple edges are allowed, but loops and cut vertices are forbidden. The outer face is allowed to have any degree. In addition, one edge on the outer face, called the *root*, is specially marked by drawing an arrow in the counterclockwise direction. Some examples are given in Figure 1.2.1.




Figure 1.2.1: Three rooted near-triangulations. From left to right, the degrees of the outer faces are 2, 3, and 5.

For each $t \in \mathcal{T}$ we define $\text{wt}_1(t)$ to be the degree of the outer face of t , and $\text{wt}_2(t)$ to be the number of interior faces of t . Then the generating function for \mathcal{T} with respect to wt_1 and wt_2 is

$$T(x, y) = \sum_{t \in \mathcal{T}} x^{\text{wt}_1(t)} y^{\text{wt}_2(t)},$$

Also, let $\mathcal{S} = \{t \in \mathcal{T} \mid \text{wt}_1(t) = 2\}$, the subset of elements of \mathcal{T} in which the outer face has degree 2. The generating function for \mathcal{S} with respect to wt_2 is

$$S(y) = \sum_{s \in \mathcal{S}} y^{\text{wt}_2(s)}.$$

Let γ be the single edge element of \mathcal{T} (which is also in \mathcal{S}). We have the simple recursive description of \mathcal{T} , which is a set bijection, given by

$$\mathcal{T} \leftrightarrow \{\gamma\} \sqcup (\{\gamma\} \times \mathcal{T} \times \mathcal{T}) \sqcup (\{\gamma\} \times (\mathcal{T} \setminus \mathcal{S})). \quad (1.2.1)$$

The second and third terms of \mathcal{T} arise as follows. If we remove the root edge from a near-triangulation $t \in \mathcal{T}$ there are two possible cases: One possibility is that we obtain two near-triangulations joined together at a cut-vertex (see Figure 1.2.2). The other possibility is that we obtain a single near-triangulation and the degree of the outer face goes up by one (see Figure 1.2.3. In this latter case, the resulting near-triangulation cannot be in \mathcal{S}).




Figure 1.2.2: Decomposition of \mathcal{T} for second term.




Figure 1.2.3: Decomposition of \mathcal{T} for third term.

Taking care to match up the weights correctly, from (1.2.1), we obtain

$$T(x, y) = x^2 + x^{-1}yT(x, y)^2 + x^{-1}y(T(x, y) - x^2S(y)) ,$$

which is a (quadratic) functional equation for $T(x, y)$ but also involves $S(y)$ linearly. To solve this equation is more complicated — the following approach is called the *quadratic method*.

First, we multiply both sides of the quadratic equation by $4xy$. This allows us to complete the square, to obtain

$$(2yT + y - x)^2 = (x - y)^2 + 4x^2y^2S - 4x^3y.$$

Let $D(x, y)$ be given by the right hand side (also the left hand side) of this equation. Now, suppose that $\alpha(y)$ is a formal power series such that $D(\alpha, y) = 0$. (For now, we're still playing fast and loose with formal power series, so we won't worry too hard about why this substitution is legal or why such an $\alpha(y)$ exists.) Then, because D is a perfect square, we also have

$$\left(\frac{\partial}{\partial x} D(x, y) \right) \Big|_{x=\alpha(y)} = 0.$$

This gives us the two equations

$$(\alpha - y)^2 + 4\alpha^2y^2S - 4\alpha^3y = 0, \quad (1.2.2)$$

$$2(\alpha - y) + 8\alpha y^2S - 12\alpha^2y = 0. \quad (1.2.3)$$

Now eliminate S between these equations, by computing 2 times (1.2.2) minus α times (1.2.3), to obtain

$$-2y(\alpha - y) + 4\alpha^3y = 0,$$

which simplifies to

$$y = \alpha(1 - 2\alpha^2). \quad (1.2.4)$$

Substituting (1.2.4) into (1.2.3) and simplifying, we obtain

$$S = \frac{1 - 3\alpha^2}{(1 - 2\alpha^2)^2}. \quad (1.2.5)$$

But now we can use LIFT, to obtain for $n \geq 1$,

$$\begin{aligned} [y^{2n}]S &= [y^{2n}] \frac{1 - 3\alpha^2}{(1 - 2\alpha^2)^2} \\ &= \frac{1}{2n} [\lambda^{2n-1}] \left(\frac{d}{d\lambda} \frac{1 - 3\lambda^2}{(1 - 2\lambda^2)^2} \right) (1 - 2\lambda^2)^{-2n} \\ &= \frac{1}{2n} [\lambda^{2n-1}] 2\lambda(1 - 6\lambda^2)(1 - 2\lambda^2)^{-2n-3} \\ &= \frac{2^{n-1}}{n} \left(\binom{3n+1}{n-1} - 3 \binom{3n}{n-2} \right) \\ &= \frac{(3n)! 2^{n+1}}{n!(2n+2)!}. \end{aligned}$$

This formula also works for $n = 0$ (it has value 1). For example, when $n = 2$, the formula has value 4, and the 4 elements of \mathcal{S} with $2 \cdot 2 = 4$ inner faces are pictured in Figure 1.2.4. Note that every element of \mathcal{S} has an even number of internal faces, since all internal faces have degree 3, and the sum of the face-degrees must be even (it equals twice the number of edges).




Figure 1.2.4: The elements of \mathcal{S} with 4 inner faces.

It is possible to deduce the coefficients in $T(x, y)$ itself, but we won't do that here; this can be found on pages 143–145 of Goulden and Jackson's "Combinatorial Enumeration". Instead we'll continue in a different direction.

Let $\mathcal{R} = \{t \in \mathcal{T} \mid \text{wt}_1(T) = 3\}$ be the subset of elements of \mathcal{T} in which the outer face has degree 3. Let $\mathcal{Q} \subset \mathcal{R}$ be the subset of rooted triangulations with no multiple edges. Consider the generating functions

$$R(y) = \sum_{t \in \mathcal{R}} y^{\text{wt}_2(t)+1} \quad \text{and} \quad Q(y) = \sum_{t \in \mathcal{Q}} y^{\text{wt}_2(t)+1},$$

with respect to the total number of faces (number of interior faces plus 1). We have a much simpler decomposition of \mathcal{S} similar to (1.2.1):

$$\mathcal{S} \leftrightarrow \{\gamma\} \sqcup \{\gamma\} \times \mathcal{R},$$

obtained by decomposing along the root edge. Thus $R(y) = S(y) - 1$.

To determine $Q(y)$, note that the following is a bijective construction for all elements of \mathcal{R} : take an element $q \in \mathcal{Q}$, and replace each edge by an element of \mathcal{S} . (Actually, there are two ways one could replace an edge by an element of \mathcal{S} . To make this replacement canonical, run your favourite algorithm on q to orient each of the edges; then insert the element of \mathcal{S} so that the root edge aligns with this orientation.) Let the root edge of the element of \mathcal{S} that replaces the root edge of q become the root edge of the element of \mathcal{R} that is constructed.

For example, in Figure 1.2.5 we show an element $A \in \mathcal{R}$ that is constructed in this way from an element $B \in \mathcal{Q}$; edges 1, 2, 3, 4 of B are replaced by $s_1, s_2, s_3, s_4 \in \mathcal{S}$, respectively, and the remaining edges of B are replaced by γ .

This construction suggests that we might expect a relationship between the generating functions $R(y)$ and $Q(y)$, via some sort of substitution involving $S(y)$. However, it is not immediately obvious what that relationship will be. To work out the details (correctly!), it is best to go back to the definition of the generating functions.

Suppose that $Q(y) = \sum_{n \geq 1} q_{2n} y^{2n}$. Each element of \mathcal{Q} with $2n$ faces has $3n$ edges. If $s_1, \dots, s_{3n} \in \mathcal{S}$ are used to replace these edges, then the resulting element of \mathcal{R} will have




Figure 1.2.5: An example of the compositional construction for \mathcal{R} .

exactly $2n + \text{wt}_2(s_1) + \dots + \text{wt}_2(s_{3n})$ faces. Thus, we have

$$\begin{aligned}
 R(y) &= \sum_{n \geq 1} q_{2n} \sum_{s_1, \dots, s_{3n} \in \mathcal{S}} y^{2n + \text{wt}_2(s_1) + \dots + \text{wt}_2(s_{3n})} \\
 &= \sum_{n \geq 1} q_{2n} y^{2n} \left(\sum_{s \in \mathcal{S}} y^{\text{wt}_2(s)} \right)^{3n} \\
 &= \sum_{n \geq 1} q_{2n} (yS(y)^{\frac{3}{2}})^{2n} \\
 &= Q(yS(y)^{\frac{3}{2}}).
 \end{aligned}$$

Hence,

$$S(y) - 1 = Q(yS(y)^{\frac{3}{2}}).$$

To solve for $Q(y)$, we write y and $S(y)$ in terms of α , using (1.2.4) and (1.2.5). This gives

$$\frac{1 - 3\alpha^2}{(1 - 2\alpha^2)^2} - 1 = Q \left(\alpha(1 - 2\alpha^2) \left(\frac{1 - 3\alpha^2}{(1 - 2\alpha^2)^2} \right)^{\frac{3}{2}} \right), \quad (1.2.6)$$

which is an equation of the form $Q(g(\alpha)) = f(\alpha)$. We've already learned how to solve this type of equation at the start of this section: $[y^n]Q(y) = [\lambda^{-1}] \frac{f'(\lambda)}{n g(\lambda)^n}$.

It turns out that the calculation works out much more cleanly if we first make the substitution

$$\alpha = \frac{\beta}{\sqrt{1 + 2\beta^2}}.$$

(This is not to say that one has to make the substitution to get the right answer, but we get a nicer formula by doing so.) Plugging this into (1.2.6) we obtain the much cleaner equation

$$\beta^2 - 2\beta^4 = Q(\beta(1 - \beta^2)^{\frac{3}{2}}).$$

Therefore for $n \geq 1$,

$$\begin{aligned} [y^{2n}]Q(y) &= [\lambda^{-1}] \frac{\frac{d}{d\lambda}(\lambda^2 - 2\lambda^4)}{2n(\lambda(1 - \lambda^2)^{\frac{3}{2}})^{2n}} \\ &= [\lambda^{2n-1}] \frac{2\lambda - 8\lambda^4}{2n(1 - \lambda^2)^{3n}} \\ &= \frac{1}{n} \binom{4n-2}{n-1} - \frac{4}{n} \binom{4n-3}{n-2} \\ &= \frac{2(4n-3)!}{n!(3n-1)!} \end{aligned}$$

(Note: In this case the formula is not valid for $n = 0$. We have $[y^0]Q(y) = 0$.)

1.3 The ring of formal power series

At this point, we've played a number of tricks with formal power series, with varying degrees of justification. However, it would be a good idea to know, for example, that the power series $\alpha(y)$ in the previous example really does exist. Otherwise, everything we've done are based on a faulty premise, and therefore nonsense. Our next goal is to put the algebra of formal power series on solid foundations, so that we know what types of operations are valid, and what types of functional equations are guaranteed to be solvable.

We begin with a review of basic ring theory. The notes here are quite sparse — any abstract algebra text will provide a more complete picture. The rings we will primarily be working with are (commutative) integral domains, and fields. I won't include all the definitions here. Briefly a commutative ring is an algebraic object in which we can add, subtract and multiply, subject to the usual rules. An integral domain has the property that the additional property that the product of non-zero elements is non-zero. Integral domains particularly nice commutative rings, as they have the cancellation property: if $a \neq 0$ and $ab = ac$ then $b = c$. A field is an integral domain in which we can divide by any non-zero element, which makes the cancellation property clear.

Some examples of integral domains are: the integers \mathbb{Z} ; polynomials with integer coefficients $\mathbb{Z}[x]$, or complex coefficients $\mathbb{C}[x]$; multivariable polynomial rings such as $\mathbb{Z}[x, y, z]$. Some examples of fields are: the rational numbers \mathbb{Q} ; the complex numbers \mathbb{C} ; fields of

rational functions such as $\mathbb{Q}(x)$, $\mathbb{C}(x)$. An example of a commutative ring that is not an integral domain is the vector space \mathbb{R}^2 , with multiplication defined by $(a, b)(c, d) = (ac, bd)$.

If R, R' are rings, a map $\phi : R \rightarrow R'$ is a *homomorphism* if $\phi(a + b) = \phi(a) + \phi(b)$, $\phi(ab) = \phi(a)\phi(b)$, and $\phi(1) = 1$. An invertible homomorphism is called an *isomorphism*. Rings which are isomorphic are essentially the same.

Proposition 1.3.1. *If R is an integral domain, then there is a unique (up to unique isomorphism) field \mathbb{F} containing R , such that for every $a \in \mathbb{F}$ there exist $p, q \in R$ such that $a = pq^{-1}$ in \mathbb{F} .*

This field \mathbb{F} is called the *field of fractions* of R . For example, the field of fractions of \mathbb{Z} is \mathbb{Q} ; the field of fractions of $\mathbb{Z}[x]$ is $\mathbb{Q}(x)$.

If $a, a', d \in R$, we write $a \equiv a' \pmod{d}$ if $a - a' = de$ for some $e \in R$. If $a \equiv a' \pmod{d}$ and $b \equiv b' \pmod{d}$ then

$$\begin{aligned} a + b &\equiv a' + b' \pmod{d} \\ ab &\equiv a'b' \pmod{d} \end{aligned}$$

Example 1.3.2. Let $R = \mathbb{Q}[x]$. Then $f(x) \equiv g(x) \pmod{x^n}$ if and only if $[x^k]f(x) = [x^k]g(x)$ for all $k < n$. In other words, doing algebra modulo x^n amounts to ignoring all terms involving x^m , $m \geq n$.

For the rest of this chapter, R will be an integral domain. We now define the ring of formal power series $R[[x]]$. (This can also be defined if R is any ring, but may not have all the properties we discuss here.)

As a set, $R[[x]]$ is the set of all expressions of the form:

$$\sum_{n \geq 0} a_n x^n,$$

where $a_0, a_1, a_2, \dots \in R$. Formally these are the same as infinite sequences of elements of R , which we've chosen to write in a suggestive way. We define addition and multiplication as follows. Put

$$A(x) = \sum_{n \geq 0} a_n x^n, \quad B(x) = \sum_{n \geq 0} b_n x^n.$$

Then

$$\begin{aligned} A(x) + B(x) &= \sum_{n \geq 0} (a_n + b_n) x^n \\ A(x)B(x) &= \sum_{n \geq 0} \left(\sum_{k=0}^n a_k b_{n-k} \right) x^n. \end{aligned}$$

These are the same formulae that are used to define addition and multiplication of polynomials.

Proposition 1.3.3. *$R[[x]]$ is an integral domain.*

To prove this is a matter of checking all axioms, which is straightforward. The fact that the product of non-zero elements is non-zero follows from considering the leading term (the term of lowest degree) of the product.

As before, we'll use the notation $[x^n]A(x) = a_n$ to denote the coefficients of x^n in $A(x)$. An important subset of $R[[x]]$ is

$$R[[x]]_+ = \{B(x) \in R[[x]] \mid [x^0]B(x) = 0\} .$$

If $B(x) \in R[[x]]_+$ we define composition:

$$A(B(x)) = \sum_{i \geq 0} a_i B(x)^i = \sum_{n \geq 0} \left(\sum_{\substack{i, j_1, \dots, j_i \geq 1 \\ j_1 + \dots + j_i = n}} a_i b_{j_1} \dots b_{j_i} \right) x^n . \quad (1.3.1)$$

It is worth taking a moment to think about why the expansion on the right hand side is correct. Note that $A(0) = a_0$.

Proposition 1.3.4. *For any $B(x) \in R[[x]]_+$, the map $\text{ev}_{B(x)} : R[[x]] \rightarrow R[[x]]$ defined by $A(x) \mapsto A(B(x))$ is a ring homomorphism. If $B(x) \neq 0$ then this map is injective. If $B(x) = 0$ then $\text{ev}_{B(x)} : R[[x]] \rightarrow R$.*

The injectivity statement is key to justifying one of the steps from the example in Section 1.2. We had an equation of the form $Q(g(\alpha)) = f(\alpha)$, where $\alpha = \alpha(y)$ was a formal power series. The way we proceeded, however, was to treat α like an indeterminate. Proposition 1.3.4 tells us that this is okay: we can undo the substitution $y \mapsto \alpha(y)$ to obtain $Q(g(y)) = f(y)$.

Another application of substitution is the construction of inverses. It is easy to see that if a_0 is an invertible element of R , then

$$(a_0 - x) \left(\sum_{n \geq 0} a_0^{-n-1} x^n \right) = 1 .$$

We can substitute $x \mapsto a_0 - A(x)$ to obtain

$$A^{-1} = (a_0 - (a_0 - A(x)))^{-1} = \sum_{n \geq 0} a_0^{-n-1} (a_0 - A(x))^n .$$

Conversely it is easy to see that if a_0 is not invertible then $A(x)$ is not invertible.

We define the formal differentiation operator

$$\frac{d}{dx} A(x) = A'(x) = \sum_{n \geq 1} n a_n x^{n-1} ,$$

and if $\mathbb{Q} \subset R$ we define its right-inverse, the formal integration operator

$$I_x A(x) = \sum_{n \geq 0} \frac{a_n}{n+1} x^{n+1} .$$

Note that if $A'(x) = B'(x)$ and $A(0) = B(0)$ then we can conclude that $A(x) = B(x)$. As we'll see later, differentiating both sides can be a powerful method of showing that two formal power series are equal.

Formal differentiation satisfies all the usual rules from calculus.

Proposition 1.3.5. *Let $A(x), B(x) \in R[[x]]$.*

(1) *For any $c \in R$, $\frac{d}{dx}(A(x) + cB(x)) = A'(x) + cB'(x)$.*

(2) *$\frac{d}{dx}(A(x)B(x)) = A'(x)B(x) + A(x)B'(x)$*

(3) *If $B(x) \in R[[x]]_+$ then $\frac{d}{dx}A(B(x)) = A'(B(x))B'(x)$.*

Proof. (1) is easy. For (2), we have

$$\begin{aligned} \frac{d}{dx}(A(x)B(x)) &= \frac{d}{dx} \sum_{n \geq 0} \left(\sum_{k=0}^n a_k b_{n-k} \right) x^n \\ &= \sum_{n \geq 0} \left(\sum_{k=0}^n (k + (n - k)) a_k b_{n-k} \right) x^{n-1} \\ &= \sum_{n \geq 0} \sum_{k=0}^n k a_k x^{k-1} b_{n-k} x^{n-k} + \sum_{n \geq 0} \sum_{k=0}^n (n - k) b_{n-k} x^{n-k-1} a_k x^k \\ &= A'(x)B(x) + A(x)B'(x). \end{aligned}$$

For (3), use (2) to prove $\frac{d}{dx}B(x)^n = nB(x)^{n-1}B'(x)$ by induction, and then

$$\frac{d}{dx}A(B(x)) = \sum_{n \geq 0} a_n \frac{d}{dx}B(x)^n = \sum_{n \geq 0} n a_n B(x)^{n-1} B'(x) = A'(B(x))B'(x). \quad \square$$

For $0 \neq A(x) \in R[[x]]$, define $\text{val } A(x)$ to be the smallest number m such that $[x^m] \neq 0$; thus $\text{val } A(x) = m$ means $A(x) = \sum_{n \geq m} a_n x^n$ with $a_m \neq 0$. Define $\text{val}(0) = +\infty$. Note that val has the following properties:

(i) $\text{val}(A(x) + B(x)) \geq \min(\text{val } A(x), \text{val } B(x))$

(ii) $\text{val}(A(x)B(x)) = \text{val } A(x) + \text{val } B(x)$.

A function on a ring with properties (i) and (ii) is called a *valuation*.

For any real number ϵ , $0 < \epsilon < 1$, define a norm $\|\cdot\|_\epsilon$ on $R[[x]]$ by $\|A(x)\|_\epsilon = \epsilon^{\text{val } A(x)}$ and $\|0\|_\epsilon = 0$. Properties (i) and (ii) imply that this norm satisfies the triangle inequality, and therefore gives $R[[x]]$ the structure of a metric space, with metric $d_\epsilon(A(x), B(x)) = \|A(x) - B(x)\|_\epsilon$. We therefore obtain a notion of a limit with respect to this metric: $\lim_{m \rightarrow \infty} A_m(x) = A(x)$ means $\lim_{m \rightarrow \infty} d_\epsilon(A_m(x), A(x)) = 0$. This is a complete metric space: it is the completion of the polynomial ring $R[x]$.

If $R = \mathbb{R}$, this notion of a limit is quite different from what one might expect. For example, consider the infinite sequence $A_m(x) = \frac{1}{2^m} + x$. You might expect $\lim A_m(x) = x$, but this is wrong, because $d_\epsilon(A_m(x), x) = 1$ for all m — in fact $A_m(x)$ is a divergent sequence. The next proposition describes what convergent sequences and series look like.

Proposition 1.3.6. *Let $A_m(x)$ be a sequence of elements of $R[[x]]$.*

- (1) $\lim_{m \rightarrow \infty} A_m(x) = A(x)$ if and only if for every $n \geq 0$ there exists $N \geq 0$ such that for all $M \geq N$, $[x^n]A_M(x) = [x^n]A(x)$. In other words, the limit exists if and only if for each n , sequence $[x^n]A_m(x)$ is eventually constant.
- (2) The series $\sum_{m \geq 0} A_m(x) = \lim_{M \rightarrow \infty} \sum_{m=0}^M A_m(x)$ converges if and only if $\lim_{m \rightarrow \infty} \text{val } A_m(x) = +\infty$ ($\iff \lim_{m \rightarrow \infty} \|A_m(x)\|_\epsilon = 0$).

In the Section 1.1 we said that we would only allow ourselves to perform operations on formal power series if any coefficient of the result could be computed by a finite algebraic process. Statement (1) tells that taking limits is a valid operation. As a result, we find that Statement (2) is much stronger than the analogous result for real numbers. For us it will be impossible to have a sequence where the norms of the terms tend to zero, but the sum is divergent.

A few of things to note: First of all, although we originally defined the expression $A(x) = \sum_{n \geq 0} a_n x^n$ as a purely formal object, we can now interpret the right hand side as an infinite sum; the sum is convergent, and it converges to the the formal object $A(x)$. Second, consider the expression we had for $A(B(x)) = \sum_{n \geq 0} a_n B(x)^n$. We can now make sense of the right hand side as an infinite series. This series is convergent if $B(x) \in R[[x]]_+$, and the limit is exactly given by the right hand side of (1.3.1). (Substitution is continuous!) Third, as one of our final steps in the proof of Proposition 1.3.5, we brought the operator $\frac{d}{dx}$ inside of an infinite sum. It is possible to justify this, but we now have the luxury of doing so simply by the fact that $\frac{d}{dx}$ is a continuous operator.

Finally, in some contexts, it can be useful to use slight variations on the definition of the valuation. For example, consider $S = R[[y]]$, where R is an integral domain. Then $S[[x]] = R[[y]][[x]] = R[[x, y]]$, the ring of formal power series in two variables. The elements of this ring are of the form

$$A(x, y) = \sum_{m, n \geq 0} a_{mn} x^m y^n, \quad a_{mn} \in R.$$

Since we'll normally want to put the variables x and y on equal footing, we define a new valuation on $R[[x, y]]$:

$$\text{val } A(x, y) = \min\{m + n \mid a_{mn} \neq 0\}.$$

This gives $R[[x, y]]$ the structure of a complete metric space; in fact it is the completion of the polynomial ring $R[x, y]$.

1.4 Hensel's Lemma

Many theorems from (single and multi-variable) real and complex analysis have analogues for Formal power series. Here's one very useful example.

Theorem 1.4.1 (Taylor's formula). *Let $A(x) \in R[[x]]$ be a formal power series. Then in $R[[u, v]]$ we have the identity*

$$A(u) = \sum_{k \geq 0} \frac{(u-v)^k}{k!} A^{(k)}(v).$$

Proof. Let $A(x) = \sum_{n \geq 0} a_n x^n$. Then

$$\begin{aligned} A(u) &= \sum_{n \geq 0} a_n ((u-v) + v)^n \\ &= \sum_{n \geq 0} \sum_{k \geq 0} a_n \binom{n}{k} (u-v)^k v^{n-k} \\ &= \sum_{k \geq 0} \frac{(u-v)^k}{k!} \sum_{n \geq 0} a_n n(n-1) \dots (n-k+1) v^{n-k} \\ &= \sum_{k \geq 0} \frac{(u-v)^k}{k!} A^{(k)}(v). \end{aligned} \quad \square$$

Corollary 1.4.2. *For any $m \geq 0$ we have,*

$$A(u) \equiv \sum_{k=0}^m \frac{(u-v)^k}{k!} A^{(k)}(v) \pmod{(u-v)^{m+1}}$$

We can now state and prove a version of Hensel's Lemma, which is an excellent tool for proving existence and uniqueness of solutions to functional equations involving formal power series. This remarkable result is actually true in a much more general context: $R[[t]]$ can be replaced by any complete valuation ring, and there is a multivariable version as well.

For the rest of this section (and only for this section) we'll use the notation $F'(t, x)$ to mean partial derivative with respect to x , i.e. $F'(t, x) = \frac{\partial}{\partial x} F(t, x)$.

Theorem 1.4.3 (Hensel's Lemma). *Let $F(t, x) \in R[[t, x]]$ be such that $F(0, 0) = 0$ and $F'(0, 0)$ is invertible. Then there exists a unique function $f(t) \in R[[t]]_+$ such that $F(t, f(t)) = 0$.*

Note that if we replace formal power series, by continuously differentiable functions defined on a neighbourhood of 0, Hensel's Lemma becomes the implicit function theorem in single variable calculus. The multivariable version of Hensel's lemma is the analogue of the implicit function theorem in multivariable calculus.

Proof. We define a sequence $f_0(t) = 0, f_1(t), f_2(t) \dots$ in $R[[t]]$, using the recursive formula:

$$f_{n+1}(t) = f_n(t) - \frac{F(t, f_n(t))}{F'(t, f_n(t))}. \quad (1.4.1)$$

By an easy inductive argument, $f_n(0) = 0$, and since $F'(0, 0)$ is invertible, this implies that $F'(t, f_n(t))$ is invertible, for all n .

Now we apply the corollary to Taylor's formula, with $A(x) = F(t, x)$, $u = f_{n+1}(t)$, $v = f_n(t)$, $m = 1$, to obtain

$$F(t, f_{n+1}(t)) \equiv F(t, f_n(t)) + (f_{n+1}(t) - f_n(t))F'(t, f_n(t)) \pmod{(f_{n+1}(t) - f_n(t))^2}.$$

From (1.4.1), the right hand side is zero; thus $(f_{n+1}(t) - f_n(t))^2$ divides $F(t, f_{n+1}(t))$, and hence $F(t, f_n(t))^2$ divides $F(t, f_{n+1}(t))$. From this we see that $\text{val} F(t, f_{n+1}(t)) \geq 2 \text{val} F(t, f_n(t)) > 0$; so

$$\lim_{n \rightarrow \infty} \text{val}(f_{n+1}(t) - f_n(t)) = \lim_{n \rightarrow \infty} \text{val} F(t, f_n(t)) = +\infty.$$

Thus we can define $f(t) = \lim_{n \rightarrow \infty} f_n(t)$. By Proposition 1.3.6, this limit exists, and since substitution is continuous,

$$F(t, f(t)) = \lim_{n \rightarrow \infty} F(t, f_n(t)) = 0.$$

To show uniqueness, suppose we had two solutions $f(t), \hat{f}(t) \in R[[t]]_+$ such that $F(t, f(t)) = F(t, \hat{f}(t)) = 0$. Then we apply the corollary to Taylor's formula with $A(x) = F(t, x)$, $u = \hat{f}(t)$, $v = f(t)$, $m = 1$, to obtain

$$F(t, \hat{f}(t)) \equiv F(t, f(t)) + (\hat{f}(t) - f(t))F'(t, f(t)) \pmod{(\hat{f}(t) - f(t))^2}.$$

Since $F'(t, f(t))$ is invertible, this gives us

$$\hat{f}(t) - f(t) \equiv 0 \pmod{(\hat{f}(t) - f(t))^2},$$

which implies that $\hat{f}(t) - f(t) = 0$. □

Note that Equation (1.4.1) is the same formula used in Newton's method for finding zeros of equations. Given a continuously differentiable function $G(x)$, and an approximate root $G(x_0) \approx 0$, Newton's method says to construct the sequence $x_{n+1} = x_n - G(x_n)/G'(x_n)$, whose limit (under reasonable hypotheses) will be a root of $G(x)$.

Example 1.4.4. As an example of Hensel's lemma, we'll prove that the binomial $1 + t$ has a unique n^{th} root in $\mathbb{C}[[t]]$ with constant term 1. Let $F(t, x) = 1 + t - (1 + x)^n$. We have $F(0, 0) = 0$, $F'(0, 0) = n \neq 0$, so there is a unique $f(t) \in \mathbb{C}[[t]]_+$ such that $F(t, f(t)) = 0$, i.e. $1 + t = (1 + f(t))^n$. Thus $1 + f(t)$ is the unique n^{th} root of $(1 + t)$ with constant term 1.

1.5 Examples of formal power series

There are a number of special formal power series that come up frequently:

$$e^x = \sum_{n \geq 0} \frac{x^n}{n!} \quad \log(1 - x)^{-1} = \sum_{n \geq 1} \frac{x^n}{n} \quad (1 + x)^a = \sum_{n \geq 0} \binom{a}{n} x^n,$$

where $\binom{a}{n} = \frac{a(a-1)\dots(a-n+1)}{n!}$. These formal power series get their names because they have many of the familiar properties of the analogous functions of a real/complex variable. However, it is worth emphasizing that when we write e^x , we will actually mean the formal power series above, not the analogous function. Here are some of the properties:

- The formal power series e^x is the unique solution to the differential equation $\frac{d}{dx}E(x) = E(x)$ with $E(0) = 1$. To prove, this, we write $E(x) = \sum \frac{e_n}{n!}x^n$, and deduce that $1 = e_0 = e_1 = e_2 = \dots$.
- $e^{x+y} = e^x e^y$, which can be proved by computing the coefficient of $x^m y^n$ on both sides. From this we deduce that $e^{-x} = 1/e^x$.
- $\frac{d}{dx} \log(1-x)^{-1} = (1-x)^{-1}$ can be checked directly.
- $\log(e^x) = x$. Before we prove this, we first need to make sense of the statement. The left hand side, is obtained by substituting $1 - e^{-x}$ into the series $\log(1-x)^{-1}$. Now to prove it we take derivatives of both sides: $\frac{d}{dx} \log(e^x) = \frac{d}{dx} \log(1 - (1 - e^{-x}))^{-1} = (1 - (1 - e^{-x}))^{-1} \frac{d}{dx}(1 - e^{-x}) = e^x e^{-x} = 1 = \frac{d}{dx} x$. Since $[x^0] \log(e^x) = 0 = [x^0] x$, the result now follows.
- $\frac{d}{dx}(1+x)^a = a(1+x)^{a-1}$ is also easy to check directly.

Exercise 1.5.1. Prove that $\log(1+x)^a = a \log(1+x)$, and that $(1+x)^a (1+x)^b = (1+x)^{a+b}$.

Other analytic functions also have formal power series analogues. Examples include:

$$\begin{aligned} \sin x &= \sum_{n \geq 0} (-1)^n \frac{x^{2n-1}}{(2n-1)!} & \cos x &= \sum_{n \geq 0} (-1)^n \frac{x^{2n}}{(2n)!} \\ \sinh x &= \sum_{n \geq 0} \frac{x^{2n-1}}{(2n-1)!} & \cosh x &= \sum_{n \geq 0} \frac{x^{2n}}{(2n)!} \\ \tan x &= \frac{\sin x}{\cos x} & \sec x &= \frac{1}{\cos x} \\ \arcsin x &= I_x((1-x^2)^{-1/2}) & \arctan x &= I_x((1+x^2)^{-1}) \end{aligned}$$

which have all the usual properties.

Example 1.5.2. As an example of how Lagrange's theorem can be used to prove identities, we prove "Abel's extension of the binomial theorem": For $n \geq 0$,

$$(\alpha + \beta)(\alpha + \beta + n)^{n-1} = \alpha \beta \sum_{k=0}^n \binom{n}{k} (\alpha + k)^{k-1} (\beta + n - k)^{n-k-1}.$$

The right hand side looks a little like the coefficient of x^n in a product of two series. To start, we rewrite the formula to make this even more so.

$$\frac{\alpha + \beta}{n!} (\alpha + \beta + n)^{n-1} = \sum_{k=0}^n \frac{\alpha}{k!} (\alpha + k)^{k-1} \frac{\beta}{(n-k)!} (\beta + n - k)^{n-k-1}.$$

Now, the identity we want to prove looks like

$$[x^n] F(\alpha + \beta, x) = [x^n] F(\alpha, x) F(\beta, x), \tag{1.5.1}$$

where

$$F(\gamma, x) = \sum_{n \geq 0} \frac{\gamma}{n!} (\gamma + n)^{n-1} x^n.$$

We need to find a more useful description of the function F . To this end, we'll try to write $[x^n]F(\gamma, x)$ as something that might appear on the right hand side of LIFT. Thus we want to find a pair of functions $\phi(\lambda)$, $f_\gamma(\lambda)$ such that

$$\frac{1}{n} [\lambda^{n-1}] f'_\gamma(\lambda) \phi(\lambda)^n = \frac{\gamma}{n!} (\gamma + n)^{n-1},$$

or rather

$$[\lambda^{n-1}] f'_\gamma(\lambda) \phi(\lambda)^n = \gamma \frac{(\gamma + n)^{n-1}}{(n-1)!} = [\lambda^{n-1}] \gamma e^{(\gamma+n)\lambda}$$

With a little trial and error, we find a nice simple pair that work: $f_\gamma(\lambda) = e^{\gamma\lambda}$ and $\phi(\lambda) = e^\lambda$.

Now, we're ready to prove the identity. We put $\phi(\lambda) = e^\lambda$ and $f_\gamma(\lambda) = e^{\gamma\lambda}$. We define $T(x)$ to be the solution to the functional equation $T(x) = x\phi(T(x)) = xe^{T(x)}$. Then we have

$$\begin{aligned} [x^n]F(\gamma, x) &= \frac{1}{n} [\lambda^{n-1}] f'_\gamma(\lambda) \phi(\lambda)^n \\ &= [x^n] f_\gamma(T(x)) \\ &= [x^n] e^{\gamma T(x)}, \end{aligned}$$

so we've found that (1.5.1) is just asserting that

$$[x^n] e^{(\alpha+\beta)T(x)} = [x^n] e^{\alpha T(x)} e^{\beta T(x)},$$

which is certainly true.

1.6 Formal Laurent series

The ring of Formal Laurent series $R((x))$ is defined to be the set of expressions of the form

$$A(x) = \sum_{k \geq N} a_n x^k$$

where $a_k \in R$ and $N \in \mathbb{Z}$. The difference between these and formal power series is that N is allowed to be negative. However, the sums are only infinite in one direction: only finitely many $[x^k]A(x)$ with $k < 0$ are allowed to be non-zero.

It is not hard to see that addition, subtraction, multiplication, differentiation, and valuation, can all be defined similarly to formal power series. For multiplicative inverses, we note that if $A(x) = \sum_{k \geq N} a_n x^k$ with a_N invertible, then $\tilde{A}(x) = x^{-N} A(x)$ is invertible in $R[[x]]$. So we can define $A(x)^{-1} = x^{-N} \tilde{A}(x)^{-1}$ in $R((x))$. (If a_n is a non-invertible non-zero element then $A(x)$ is not invertible in $R((x))$.) In particular, if \mathbb{F} is a field then $\mathbb{F}((x))$ is also a field. If \mathbb{F} is the field of fractions of R then $\mathbb{F}((x))$ is the field of fractions of $R[[x]]$.

If $B(x) \in R[[x]]_+$ the substitution $A(B(x))$ still makes sense, provided $B(x)$ is invertible in $R((x))$ (both in terms of each coefficient being finitely computable, and in terms of the metric space structure defined by the valuation). $A(0)$, however, is not defined when $\text{val } A(x) < 0$.

One thing that is missing from $R((x))$ is the formal integration operator. We can't define such an operator, because there is no formal Laurent series $A(x)$ such that $A'(x) = \frac{1}{x}$. In fact it is not hard to see:

Proposition 1.6.1. *For any $A(x) \in R((x))$, $[x^{-1}]A'(x) = 0$.*

Instead, we define the *formal residue* of $A(x)$ to be $[x^{-1}]A(x)$. The R -linear map $[x^{-1}]$ is called the *formal residue operator*, and it behaves like a definite integral. For example, we have an integration by parts formula:

Proposition 1.6.2. *For any $f(x), g(x) \in R((x))$ we have $[x^{-1}]f'(x)g(x) = -[x^{-1}]f(x)g'(x)$.*

To prove this, just expand $[x^{-1}]\frac{d}{dx}(f(x)g(x)) = 0$. We also have a substitution formula.

Theorem 1.6.3. *If $A(x) \in R((x))$ and $B(y) \in R[[y]]_+$ is invertible in $R((y))$, then*

$$[x^{-1}]A(x) = \frac{1}{m}[y^{-1}]A(B(y))B'(y)$$

where $m = \text{val } B(y)$.

Proof. By linearity of the formal residue operator, it is enough to prove this for $A(x) = x^n$. If $n \neq -1$ we have

$$\begin{aligned} [x^{-1}]x^n &= 0 \\ \frac{1}{m}[y^{-1}]B(y)^n B'(y) &= \frac{1}{m(n+1)}[y^{-1}]\frac{d}{dy}B(y)^{n+1} = 0 \end{aligned}$$

If $n = -1$, write $B(y) = \sum_{k \geq m} b_k y^k$. Then

$$\begin{aligned} [x^{-1}]x^{-1} &= 1 \\ \frac{1}{m}[y^{-1}]B(y)^{-1}B'(y) &= \frac{1}{m}[y^{-1}]\frac{mb_m y^{m-1} + \dots}{b_m y^m + \dots} = 1 \end{aligned}$$

□

We can now state and prove a stronger version of LIFT.

Theorem 1.6.4 (LIFT). *Let $\phi(\lambda) \in R[[\lambda]]$ be a formal power series, with $\phi(0)$ invertible. There exists a unique formal power series $A(x) \in R[[x]]_+$ satisfying*

$$A(x) = x\phi(A(x)).$$

Moreover, if $f(\lambda) \in R((\lambda))$ is any formal Laurent series, then

- (i) for $n \neq 0$, $[x^n]f(A(x)) = \frac{1}{n}[\lambda^{n-1}]f'(\lambda)\phi(\lambda)^n$;
- (ii) $[x^0]f(A(x)) = [\lambda^0]f(\lambda) + [\lambda^{-1}]f'(\lambda) \log\left(\frac{\phi(\lambda)}{\phi(0)}\right)$.

Proof. To prove the existence and uniqueness of $A(x)$, simply apply Hensel's lemma to $F(x, y) = y - x\phi(y)$. Since $F(0, 0) = 0$, $\frac{d}{dy}F(x, y)|_{(0,0)} = 1$, there is a unique $A(x)$ such that $F(x, A(x)) = 0$.

For (i) we apply Theorem 1.6.3 with $B(\lambda) = \lambda/\phi(\lambda)$. Note that $B(A(x)) = x$, so $A(B(\lambda)) = \lambda$ and $\text{val } B(\lambda) = 1$.

$$\begin{aligned} [x^n]f(A(x)) &= [x^{-1}]x^{-n-1}f(A(x)) \\ &= [\lambda^{-1}]B(\lambda)^{-n-1}f(A(B(\lambda)))B'(\lambda) \\ &= [\lambda^{-1}]B(\lambda)^{-n-1}B'(\lambda)f(\lambda) \\ &= [\lambda^{-1}]\frac{d}{d\lambda}\left(\frac{B(\lambda)^{-n}}{-n}\right)f(\lambda). \end{aligned}$$

Now we apply Proposition 1.6.2 and continue

$$\begin{aligned} &= -[\lambda^{-1}]\frac{B(\lambda)^{-n}}{-n}f'(\lambda) \\ &= \frac{1}{n}[\lambda^{-1}]\phi(\lambda)^n\lambda^{-n}f'(\lambda) \\ &= \frac{1}{n}[\lambda^{n-1}]f'(\lambda)\phi(\lambda)^n. \end{aligned}$$

The proof of (ii) is similar; we leave it as exercise. □

Chapter 2

Exponential generating functions

Exponential generating functions are designed for a slightly different purpose than ordinary generating functions; namely they are used to count the number of combinatorial structures of a given type on any finite set X . Some examples of combinatorial structures on a finite set are trees and graphs (with a given vertex set), set partitions, etc. For example, exponential generating functions are great for answering questions such as, “How many trees are there with vertex set X ?”

In this chapter, we introduce exponential generating functions through the formalism of “species”. We’ll also discuss mixed (part exponential, part ordinary) generating functions and as an application, sketch a combinatorial proof of Lagrange’s implicit function theorem.

2.1 Species

We begin by introducing some general terminology to talk about the different types of combinatorial structures we can have on a finite set X . We use the term “species” to refer to a type of combinatorial structure.

Definition 2.1.1. A *species* \mathcal{A} is a rule, which assigns

- (a) to each finite set X , a finite set \mathcal{A}_X of “combinatorial structures on X ”;
- (b) to each bijection of finite sets $f : X \rightarrow Y$, a bijection $f_* : \mathcal{A}_X \rightarrow \mathcal{A}_Y$;


such that

- (I) if $\text{id} : X \rightarrow X$ is the identity map, then $\text{id}_* : \mathcal{A}_X \rightarrow \mathcal{A}_X$ is the identity map;
- (II) if $f : X \rightarrow Y$ and $g : Y \rightarrow X$ are bijections, then $(g \circ f)_* = g_* \circ f_* : \mathcal{A}_X \rightarrow \mathcal{A}_Z$.


The set \mathcal{A}_X is called the set of \mathcal{A} -structures on X ; The map $f_* : \mathcal{A}_X \rightarrow \mathcal{A}_Y$ is called the transportation of \mathcal{A} -structures along f .

This might look quite involved, but a few examples should convince you that this really isn’t all that complicated.

Example 2.1.2. Let \mathcal{T} be the species of trees. For each finite set X , \mathcal{T}_X is the set of trees with vertex set X . If $f : X \rightarrow Y$ is a bijection, then f_* is the map that relabels each vertex i with $f(i)$. For example:



If $f : \{1, 2, 3, 4, 5\} \rightarrow \{11, 12, 13, 14, 15\}$ is the map $x \mapsto 16 - x$ then



As we'll see, for just about every example of a species \mathcal{A} one might wish to consider, there is a convenient pictorial way to represent the \mathcal{A} -structures on X . In this representation, certain objects in the picture will be labelled by the elements of X . If $f : X \rightarrow Y$ is a bijection, then f_* corresponds to relabelling using f . In this situation, conditions (I) and (II) are always satisfied.

Example 2.1.3. Here are some other standard examples of species:

- The species “Set”, denoted \mathcal{E} , assigns to each finite set X the singleton $\mathcal{E}_X = \{X\}$. (You should think of this as the set of all ways to put a set-structure on X . But X is already a set, so there's only one way—just leave it alone. This may sound silly now, but when we start to do complicated constructions involving various species, it's very useful to have “and then leave the set alone” as one of the options.) If $f : X \rightarrow Y$ is a bijection, then f_* is defined in the only possible way: $f_*(X) = Y$. Pictorially, the set X can be visualized as a collection of dots, labelled with the elements of X .
- The species “ m -Set”, denoted \mathcal{E}_m , assigns to a finite set X ,

$$(\mathcal{E}_m)_X = \begin{cases} \{X\} & \text{if } \#X = m, \\ \emptyset & \text{otherwise.} \end{cases}$$

(Think of this as all ways to turn X into an m -element set. There's one way if X has m elements, and it's impossible otherwise.) If $f : X \rightarrow Y$ is a bijection, then $f_*(X) = Y$ if $\#X = \#Y = m$.

- The species “Permutations”, denoted \mathcal{S} , assigns to X the set $\mathcal{S}_X = \{\pi : X \rightarrow X \mid \pi \text{ is a bijection}\}$. If $f : X \rightarrow Y$ is a bijection, then $f_*(\pi) = f \circ \pi \circ f^{-1}$. A permutation π can be visualized as a directed graph with vertex set X , where we have arrows from x to $\pi(x)$.
- The species “Cyclic Permutations”, denoted \mathcal{C} , is defined similarly. We put $\mathcal{C}_X = \{\pi : X \rightarrow X \mid \pi \text{ is a cyclic permutation}\}$. The visualization and definition of f_* are also similar to \mathcal{S} .

- The species “Linear Orders”, denoted \mathcal{L} , assigns to X the set of all ways to turn X into an ordered list. Thus $\mathcal{L}_X = \{(x_1, x_2, \dots, x_n) \mid n = \#X, \{x_1, \dots, x_n\} = X\}$. The transport of linear structures along f is given by $f_*(x_1, \dots, x_n) = (f(x_1), \dots, f(x_n))$. A linear order (x_1, \dots, x_n) can be visualized as a directed graph with arrows from x_i to x_{i+1} .

Other examples of species include graphs, connected graphs, rooted trees, endofunctions (maps from a set to itself), set partitions, etc.

If \mathcal{A} is a species, then $\#(\mathcal{A}_X)$ depends only on the size of X . If $\#X = n$ we write $\#\mathcal{A}_n = \#(\mathcal{A}_X)$. The *exponential generating function* of the species \mathcal{A} is defined to be

$$A(x) = \sum_{n \geq 0} \#\mathcal{A}_n \frac{x^n}{n!}.$$

Example 2.1.4. For example the generating functions for the species in Example 2.1.3 are:

- $\mathcal{E} : E(x) = \sum_{n \geq 0} 1 \frac{x^n}{n!} = e^x$.
- $\mathcal{E}_m : E_m(x) = \sum_{n \geq 0} \delta_{nm} \frac{x^n}{n!} = \frac{x^m}{m!}$.
- $\mathcal{S} : S(x) = \sum_{n \geq 0} n! \frac{x^n}{n!} = (1 - x)^{-1}$.
- $\mathcal{C} : C(x) = \sum_{n \geq 0} (n - 1)! \frac{x^n}{n!} = \log(1 - x)^{-1}$.
- $\mathcal{L} : L(x) = \sum_{n \geq 0} n! \frac{x^n}{n!} = (1 - x)^{-1}$.

Definition 2.1.5. A *natural transformation* of species $\tau : \mathcal{A} \rightarrow \mathcal{B}$ is a rule which assigns to each finite set X a function $\tau_X : \mathcal{A}_X \rightarrow \mathcal{B}_X$, in such a way that for any bijection $f : X \rightarrow Y$, we have $f_* \circ \tau_X = \tau_Y \circ f_*$.

If each τ_X is a bijection, we say that τ_X is a *natural equivalence* of species. We say \mathcal{A} and \mathcal{B} are *naturally equivalent* and write $\mathcal{A} \simeq \mathcal{B}$.

Naturally equivalent structures are “fundamentally the same”. Though formally they may be different, they are the same in that they can be visualized by the same picture.

Example 2.1.6. For any species \mathcal{A} , define $\{1\} \times \mathcal{A}$ be the species whose structures are $(\{1\} \times \mathcal{A})_X = \{1\} \times \mathcal{A}_X$. The map $\tau : \mathcal{A} \rightarrow \{1\} \times \mathcal{A}$ which sends an \mathcal{A} -structure α to the pair $(1, \alpha)$ is a natural equivalence. The difference between α and $(1, \alpha)$ is superficial, because the “1” is extraneous information.

Clearly naturally equivalent species will have the same generating function. This begs the obvious question: if two species have the same generating function, are they naturally equivalent? The answer is no, or we wouldn’t have given such a convoluted definition. For example, the species \mathcal{L} and \mathcal{S} are not naturally equivalent. Intuitively this should be clear: we can’t use the same types of pictures to represent a permutation and a linear order. Here is a more formal proof.

Proof. Consider $X = \{1, 2\}$ and let $f : X \rightarrow X$ be the map $f(1) = 2, f(2) = 1$. Then

$$\mathcal{L}_X = \{(1, 2), (2, 1)\} \quad \text{and} \quad f_*(a, b) = (b, a).$$

On the other hand

$$\mathcal{S}_X = \{\text{id}, f\} \quad \text{and} \quad f_*(\pi) = \pi.$$

Thus we see that if $\tau : \mathcal{L} \rightarrow \mathcal{S}$ is any natural transformation, then

$$\tau_X(1, 2) = f_*(\tau_X(1, 2)) = \tau_X(f_*(1, 2)) = \tau_X(2, 1).$$

Hence τ is not a bijection, so it can't be a natural equivalence. □

2.2 Interspecies interactions

There are several basic operations on species:

1. Sum: if $\mathcal{A}_1, \mathcal{A}_2, \dots$ is a finite (or infinite) list of species, we define

$$(\mathcal{A}_1 \oplus \mathcal{A}_2 \oplus \dots)_X = (\{1\} \times \mathcal{A}_1)_X \sqcup (\{2\} \times \mathcal{A}_2)_X \sqcup \dots$$

We put $\{i\} \times \mathcal{A}_i$ on the right hand side (rather than \mathcal{A}_i) to make sure that the union is disjoint. (If the list is infinite, this is only defined if the union is a finite set.)

Example 2.2.1. There is a natural equivalence of species $\mathcal{E} \simeq \mathcal{E}_0 \oplus \mathcal{E}_1 \oplus \mathcal{E}_2 \oplus \dots$.

2. Difference: If $\mathcal{A}_X \subset \mathcal{B}_X$ for every set X , then we define

$$(\mathcal{A} \setminus \mathcal{B})_X = \mathcal{A}_X \setminus \mathcal{B}_X.$$

3. Product: If \mathcal{A}, \mathcal{B} are species, then an $(\mathcal{A} * \mathcal{B})$ -structure on X is a pair, consisting of an \mathcal{A} -structure on S and a \mathcal{B} -structure on $X \setminus S$, where $S \subset X$. In other words:

$$(\mathcal{A} * \mathcal{B})_X = \coprod_{S \subset X} \mathcal{A}_S \times \mathcal{B}_{X \setminus S}.$$

Example 2.2.2. Consider $\mathcal{T} \times \mathcal{T}$, where \mathcal{T} is the species of trees. Then elements $(\mathcal{T} \times \mathcal{T})_X$ are ordered pair of trees (t_1, t_2) , such $X = \text{vert}(t_1) \sqcup \text{vert}(t_2)$.

For a second example, consider $\mathcal{E}_1 \times \mathcal{T}$. Then $(\mathcal{E}_1 \times \mathcal{T})_X$ consists of all pairs (x, t) where $x \in X$ and t is a tree with vertex set $X \setminus \{x\}$.

4. Marking an element: If \mathcal{A} is a species, then \mathcal{A}^\bullet -structure on X is a pair (α, x) where α is an \mathcal{A} -structure on X and $x \in X$ is a “marked element”. Thus,

$$\mathcal{A}_X^\bullet = \mathcal{A}_X \times X.$$

Pictorially, we represent the marked element by drawing some sort of special mark around it (e.g. an extra circle, a red x, etc.).

Example 2.2.3. \mathcal{T}^\bullet is the species of rooted trees. The root of the tree is the marked element.

5. Mark and delete: If \mathcal{A} is a species define

$$\mathcal{A}'_X = \mathcal{A}_{X \cup \{X\}}.$$

The idea here is to add a “special element” to the set X (in standard set theory, a set cannot contain itself, so X itself is a good choice for this element), then look at \mathcal{A} -structures on the bigger set. Pictorially, the special element is distinguished, because it’s labelled by the set X rather than an element of X ; we normally represent this with by drawing a special mark of some sort, instead of the label X .

The name “mark and delete” is a bit misleading. It comes from the fact that if you delete the label on the marked element x in the construction of \mathcal{A}'_X , you essentially get an \mathcal{A} -structure on $X \setminus \{x\}$.

Example 2.2.4. \mathcal{T}' is the species of rooted trees, where the root vertex is unlabelled, and all other vertices are labelled.

6. Composition: This one is rather involved. We need several definitions.

Definition 2.2.5. We say that a species \mathcal{A} is *connected* if $\mathcal{A}_\emptyset = \emptyset$.

Definition 2.2.6. We say that P is a *set partition* of a finite set X , and write $P \vdash X$ if $P = \{S_1, S_2, \dots, S_k\}$ where S_1, \dots, S_k are disjoint non-empty sets and $X = S_1 \sqcup \dots \sqcup S_k$.

Definition 2.2.7. If P is a set of finite sets and \mathcal{A} is a species, an \mathcal{A} -*function* on P is a function $a : P \rightarrow \mathbf{Sets}$ such that $a(S) \in \mathcal{A}_S$ for all $S \in P$.

Now we can define species composition. If \mathcal{A}, \mathcal{B} are species, and \mathcal{B} is connected, an $\mathcal{A}[\mathcal{B}]$ -structure on X consists of a set partition P of X , an \mathcal{A} -structure on P , and \mathcal{B} -function on P . Thus,

$$\mathcal{A}[\mathcal{B}]_X = \coprod_{P \vdash X} \mathcal{A}_P \times \{\mathcal{B}\text{-functions on } P\}.$$

Pictorially, we would represent an $\mathcal{A}[\mathcal{B}]$ -structure as follows. We draw an \mathcal{A} structure with a few of big circles C_1, \dots, C_k . The big circles replace the objects that normally be labelled, if we were *just* drawing a \mathcal{A} structure. But now, instead of labelling the circles, we draw a \mathcal{B} -structure on a set S_i inside each circle C_i , where the sets $\{S_1, \dots, S_k\}$ form a set partition of X .

Example 2.2.8. Figure 2.2.1 shows a picture of a $\mathcal{T}^\bullet[\mathcal{T}]$ -structure on $[10]$. A rooted tree is drawn with large circles for vertices. Then we draw a labelled tree inside each circle.

For a second example, consider $\mathcal{E}[\mathcal{T}]$. A $\mathcal{E}[\mathcal{T}]$ -structure on X consists of a set partition P (with no additional structure coming from \mathcal{E}), and a tree structure on each element of P . In other words, a $\mathcal{E}[\mathcal{T}]$ -structure on X is a forest with vertex set X .




Figure 2.2.1: A $\mathcal{T}^\bullet[\mathcal{T}]$ -structure on $[10]$.

A third example is the species $\mathcal{E}[\mathcal{C}]$. A $\mathcal{E}[\mathcal{C}]$ -structure on X is a collection of cyclic permutations $\pi_i : S_i \rightarrow S_i$, where $\{S_1, \dots, S_k\} \vdash X$. In other words, this is a permutation of X . From this, is not hard to see that we have a natural equivalence $\mathcal{E}[\mathcal{C}] \simeq \mathcal{S}$.

Theorem 2.2.9 (Main Theorem for Exponential Generating Functions). *Let \mathcal{A} and \mathcal{B} be species, with exponential generating functions $A(x)$ and $B(x)$ respectively. Then the exponential generating functions for ...*

- (i) ... $\mathcal{A} \oplus \mathcal{B}$ is $A(x) + B(x)$;
- (ii) ... $\mathcal{A} \setminus \mathcal{B}$ is $A(x) - B(x)$, if $\mathcal{A}_X \subset \mathcal{B}_X$ for all X ;
- (iii) ... $\mathcal{A} * \mathcal{B}$ is $A(x)B(x)$;
- (iv) ... \mathcal{A}^\bullet is $x \frac{d}{dx} A(x)$;
- (v) ... \mathcal{A}' is $\frac{d}{dx} A(x)$;
- (vi) ... $\mathcal{A}[\mathcal{B}]$ is $A(B(x))$, if \mathcal{B} is connected.

For example, the fact $e^{\log(1-x)^{-1}} = (1-x)^{-1}$ can be attributed to the fact that we have a natural equivalence $\mathcal{E}[\mathcal{C}] \simeq \mathcal{S}$.

Proof. We'll prove (vi) and leave the others as an exercise. We have

$$\sum_{n \geq 0} \# \mathcal{A}[\mathcal{B}]_n \frac{x^n}{n!} = \sum_{n \geq 0} \sum_{P \vdash X} \#(\mathcal{A}_P) \prod_{i=1}^k \#(\mathcal{B}_{S_i}) \frac{x^n}{n!},$$

where $P = \{S_1, \dots, S_k\}$. Now, every set partition with k elements can be ordered in $k!$ ways. Thus the right hand side equals

$$= \sum_{n \geq 0} \sum_{k \geq 1} \frac{1}{k!} \# \mathcal{A}_k \sum_{(S_1, \dots, S_k) \vdash X} \prod_{i=1}^k \#(\mathcal{B}_{S_i}) \frac{x^n}{n!}.$$




Figure 2.2.2: A tree (left) and its decomposition (right).

But there are $\frac{n!}{j_1! \dots j_k!}$ set compositions with $\#S_i = j_i$, so we can rewrite this as

$$\begin{aligned}
&= \sum_{n \geq 0} \sum_{k \geq 1} \frac{1}{k!} \#\mathcal{A}_k \sum_{\substack{j_1, \dots, j_k \geq 1 \\ j_1 + \dots + j_k = n}} \frac{n!}{j_1! \dots j_k!} \prod_{i=1}^k \#\mathcal{B}_{j_i} \frac{x^n}{n!} \\
&= \sum_{n \geq 0} \sum_{k \geq 1} \sum_{\substack{j_1, \dots, j_k \geq 1 \\ j_1 + \dots + j_k = n}} \frac{\#\mathcal{A}_k}{k!} \frac{\#\mathcal{B}_{j_1}}{j_1!} \dots \frac{\#\mathcal{B}_{j_k}}{j_k!} x^n \\
&= A(B(x)) \quad \square
\end{aligned}$$

Example 2.2.10. Let \mathcal{G} be the species of graphs, and let $\tilde{\mathcal{G}}$ be the species of connected graphs. The exponential generating function for \mathcal{G} is

$$G(x) = \sum_{n \geq 0} 2^{\binom{n}{2}} \frac{x^n}{n!}.$$

We have a natural equivalence $\mathcal{G} \simeq \mathcal{E}[\tilde{\mathcal{G}}]$. Thus the exponential generating function \tilde{G} for $\tilde{\mathcal{G}}$ satisfies $e^{\tilde{G}(x)} = G(x)$, and so

$$\tilde{G}(x) = \log \sum_{n \geq 0} 2^{\binom{n}{2}} \frac{x^n}{n!}.$$

Example 2.2.11. Let $T(x)$ be the exponential generating function for \mathcal{T}^\bullet , the species of rooted labelled trees. Every tree decomposes as the root, plus a set of rooted labelled trees (see Figure 2.2.2). Thus we have a natural equivalence

$$\mathcal{T}^\bullet = \mathcal{E}_1 * \mathcal{E}[\mathcal{T}^\bullet],$$

which implies

$$T(x) = x e^{T(x)}.$$

We can solve this by LIFT, to obtain

$$[x^n]T(x) = \frac{1}{n} [\lambda^{n-1}] e^{\lambda^n} = \frac{1}{n} \frac{n^{n-1}}{(n-1)!}.$$

Thus the number of rooted labelled trees with n vertices is $\#\mathcal{T}_n^\bullet = n^{n-1}$. From the relation between \mathcal{T} and \mathcal{T}^\bullet , we deduce that the number of (unrooted) labelled trees is $\#\mathcal{T}_n = n^{n-2}$.

2.3 Mixed generating functions

Definition 2.3.1. A \mathbb{Z}^k -valued *weight function* wt on a species \mathcal{A} , is a rule that assigns to each finite set X a function $\text{wt}_X : \mathcal{A}_X \rightarrow \mathbb{Z}^k$, such that for any bijection $f : X \rightarrow Y$, $\text{wt}_X = \text{wt}_Y \circ f_*$.

If \mathcal{A} and \mathcal{B} are naturally equivalent species with weight functions $\text{wt}^{\mathcal{A}}$ and $\text{wt}^{\mathcal{B}}$, we say a natural equivalence $\tau : \mathcal{A} \rightarrow \mathcal{B}$ is a *weighted natural equivalence* if for every finite set X , $\text{wt}_X^{\mathcal{A}} = \text{wt}_X^{\mathcal{B}} \circ \tau_X$.

In other words, the weight of a combinatorial object shouldn't change if we change the labellings, or under a weighted natural equivalence.

Given a species \mathcal{A} with a weight function wt , the *mixed generating function* for \mathcal{A} with respect to wt is:

$$A(x; t_1, \dots, t_k) = \sum_{n \geq 0} A_n(t_1, \dots, t_k) \frac{x^n}{n!},$$

where $A_n(t_1, \dots, t_k)$ is the ordinary generating function for $\mathcal{A}_{[n]}$ with respect to $\text{wt}_{[n]}$.

If $\text{wt}^{\mathcal{A}}$ and $\text{wt}^{\mathcal{B}}$ are weight functions on \mathcal{A} , \mathcal{B} respectively, we define weight functions on $\mathcal{A} \oplus \mathcal{B}$, $\mathcal{A} * \mathcal{B}$, \mathcal{A}^\bullet , \mathcal{A}' , and $\mathcal{A}[\mathcal{B}]$ as follows:

1. For $\gamma \in (\mathcal{A} \oplus \mathcal{B})_X$,

$$\text{wt}_X(\gamma) = \begin{cases} \text{wt}_X^{\mathcal{A}}(\gamma) & \text{if } \gamma \in \mathcal{A}_X \\ \text{wt}_X^{\mathcal{B}}(\gamma) & \text{if } \gamma \in \mathcal{B}_X \end{cases}$$

2. For $(\alpha, \beta) \in (\mathcal{A} * \mathcal{B})_X$, $\text{wt}_X(\alpha, \beta) = \text{wt}_X^{\mathcal{A}}(\alpha) + \text{wt}_X^{\mathcal{B}}(\beta)$.

3. For $(\alpha, x) \in \mathcal{A}_X^\bullet$, $\text{wt}_X(\alpha, x) = \text{wt}_X^{\mathcal{A}}(\alpha)$

4. For $\alpha \in \mathcal{A}'_X$, $\text{wt}_X(\alpha) = \text{wt}_{X \cup \{X\}}^{\mathcal{A}}(\alpha)$

5. For $(\alpha, b) \in \mathcal{A}[\mathcal{B}]$, coming from set partition $P = \{S_1, \dots, S_k\}$,

$$\text{wt}_X(\alpha, b) = \text{wt}_P^{\mathcal{A}}(\alpha) + \sum_{i=1}^k \text{wt}_{S_i}^{\mathcal{B}}(b(S_i)).$$

In each case the weight of the new object is just the sum of the weights of the sub-objects used in its construction.

Theorem 2.3.2 (Main Theorem for Mixed Generating Functions). *With the weight functions defined as above, the Main Theorem for Exponential Generating Functions is true for mixed generating functions.*

We omit the proof, as it uses essentially the same ideas as the proof of the Main Theorem for exponential generating functions.

Example 2.3.3. Define a weight function the \mathcal{C} , the species of cyclic permutations, by $\text{wt}_X(\pi) = 1$ for all $\pi \in \mathcal{C}_X$. Then $\mathcal{S} \simeq \mathcal{E}[\mathcal{C}]$ inherits the weight function

$$\text{wt}_X(\pi) = \text{number of cycles in } \pi, \quad \pi \in \mathcal{S}_X.$$

The mixed generating function for \mathcal{S} with respect to wt is

$$S(x; y) = e^{y \log(1-x)^{-1}} = (1-x)^{-y}.$$

The numbers $S_{n,k} = n![x^n y^k]S(x, y)$ are called Stirling numbers of the first kind. (The numbers $\tilde{S}_{n,k} = (-1)^{n-k} S_{n,k} = n![x^n y^k]S(-x, -y)$ are also of interest, and are called the signed Stirling numbers of the first kind.)

From $S(x; y)$ we can compute (among other things) the expected number of cycles in a permutation. If we write $S(x; y) = \sum_{n,k \geq 0} S_{n,k} y^k \frac{x^n}{n!}$, this is

$$\begin{aligned} \sum_{k \geq 0} \frac{k S_{n,k}}{n!} &= [x^n] \frac{d}{dy} S(x; y) \Big|_{y=1} \\ &= [x^n] (1-x)^{-y} \log(1-x)^{-1} \Big|_{y=1} \\ &= [x^n] (1-x)^{-1} \log(1-x)^{-1} \\ &= 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{n} \end{aligned}$$

Example 2.3.4. *another example needed*

As another application of this theory, we sketch a combinatorial proof of LIFT (the version stated in Theorem 1.1.2).

Combinatorial Proof of LIFT (sketch). We'll consider the now familiar species \mathcal{E} and \mathcal{T}^\bullet , but with some fancy new \mathbb{Z}^∞ -valued weight functions. To emphasize these special weight functions, we'll give the species new names.

Let Φ be the species \mathcal{E} with weight function

$$X \mapsto (d_0, d_1, d_2, d_3, \dots)$$

where $d_i = 0$ if $\#X \neq i$, and $d_{\#X} = 1$. Φ has mixed generating function

$$\phi(x) = \phi(x; t_0, t_1, \dots) = \sum_{n \geq 0} t_n \frac{x^n}{n!}.$$

By assigning different values to t_0, t_1, \dots , this can be any formal power series.

Define \mathcal{A} to be the species \mathcal{T}^\bullet with weight function

$$t \mapsto (c_0, c_1, c_2, c_3, \dots)$$

where c_i is the number of vertices in t rooted labelled tree t with up-degree i . For the mixed generating function $A(x) = A(x; t_0, t_1, \dots)$ for \mathcal{A} , we have chosen our weight functions so

that the decomposition of a rooted labelled tree in Example 2.2.11 is a weighted natural equivalence:

$$\mathcal{A} \simeq \mathcal{E}_1 * \Phi[\mathcal{A}].$$

Thus $A(x) = x\phi(A(x))$.

We prove LIFT, by showing that for all $m, n \geq 0$, the ordinary generating function for $\mathcal{E}_m[\mathcal{A}]_{[n]}^\bullet$, which is

$$n![x^n]x \frac{d}{dx} \frac{(A(x))^m}{m!} = \frac{n!}{m!} (n[x^n]A(x)^m),$$

equals the ordinary generating function for $(\mathcal{E}_m^\bullet * \Phi * \dots * \Phi)_{[n]}$, which is

$$n![x^n]m \frac{x^m}{m!} \phi(x)^m = \frac{n!}{m!} ([x^{n-1}] (\frac{d}{dx} x^m) \phi(x)^n).$$

To do this, we construct a weight preserving bijection

$$\begin{aligned} \mathcal{E}_m[\mathcal{A}]_{[n]}^\bullet &\rightarrow (\mathcal{E}_m^\bullet * \Phi * \dots * \Phi)_{[n]} \\ (w, i) &\mapsto (R, r, B_j, \dots, B_n, B_1, \dots, B_{j-1}) \end{aligned}$$

where

- w is a forest consisting of m rooted trees with vertices labelled
- $i \in [n]$ is one of the vertices in the forest.
- $(R, B_j, \dots, B_n, B_1, \dots, B_{j-1})$ is a “set composition” of $[n]$, where R has exactly m elements and the B_i ’s may be empty.
- $r \in R$ is a marked element in R .

Here is the idea of the bijection. Given (w, i) as above, we need to define R, r, j , and B_1, \dots, B_n . R will be the set of roots of the trees in w . Choose $r \in R$ to be the root on the same tree as i . Choose j so that i is the j^{th} smallest index on its tree. Finally B_1, \dots, B_n are the sets of children of the vertices in w —the order is obtained by running your favourite search algorithm on each tree in w , starting with r , and when all vertices of its tree are exhausted, proceeding through R in ascending cyclic order.

For example, if w is the forest shown in Figure 2.3.1, and $i = 3$, then $R = \{2, 4, 5\}$, $r = 5$ (since i is on the same component as 5), and $j = 2$ (since i is the second smallest label on its tree). Now suppose our favourite search algorithm is Breadth First Search, with vertices added to the queue in ascending order of label. We start by running this algorithm on the tree containing 5. This gives

$$B_1 = \{1, 3\} \quad B_2 = \{8\} \quad B_3 = \emptyset \quad B_4 = \emptyset,$$

at which point we’ve listed the children of every vertex on that tree. We move on to 2, which is next in R in ascending cyclic order.

$$B_5 = \{7\} \quad B_6 = \{6\} \quad B_3 = \emptyset.$$




Figure 2.3.1: A forest of rooted labelled trees.

Finally we move on to 4, which has no children, so

$$B_8 = \emptyset.$$

Thus

$$(i, w) \mapsto (5, \{2, 4, 5\}, \{8\}, \emptyset, \emptyset, \{7\}, \{6\}, \emptyset, \emptyset, \{1, 3\}).$$

The key to reversing the map is to determine j . Once j is known, we know B_1, B_2, \dots, B_n , and we can reconstruct the trees in w from this sequence because we know the algorithm that produced it. The procedure to determine j is none other than good old fashioned guess and check. First we try to see if $j = 1$ is plausible, but if that were the case, we'd have

$$B_1 = \{8\}, B_2 = \emptyset, B_3 = \emptyset, B_4 = \{7\}, B_5 = \{6\}, B_6 = \emptyset, B_7 = \emptyset, B_8 = \{1, 3\},$$

and there is no forest with 3 trees that produces this sequence under breadth first search. So $j \neq 1$. Then we try $j = 2$, which of course does work. In general j will be the smallest number for which the sequence is plausible.

This works in general, and the not-entirely-trivial details about why are left as an exercise. (You should first try to convince yourself of this when $m = 1$.) \square

Chapter 3

Signed formulae

The answer to any enumeration problem is a positive integer. However, there are a number of enumeration problems where the best known formula for this this answer is not *obviously* positive — for example, the answer may be given as the determinant of a matrix. In these situations the connection between the combinatorics and algebra becomes less transparent. A variety of techniques can be used to show that these formulae are nevertheless correct. We'll take a look at some of the most common ones in this chapter.

3.1 Inclusion-Exclusion

Theorem 3.1.1. *Let f, g be functions from the set of all subsets $\alpha \subseteq \{1, \dots, n\}$ to an abelian group (usually \mathbb{Z} or \mathbb{Q}). Then*

$$f(\alpha) = \sum_{\alpha \subseteq \beta} g(\beta), \quad \alpha \subseteq \{1, \dots, n\} \quad (3.1.1)$$

if and only if

$$g(\alpha) = \sum_{\alpha \subseteq \beta} (-1)^{|\beta| - |\alpha|} f(\beta), \quad \alpha \subseteq \{1, \dots, n\} \quad (3.1.2)$$

Proof. It is not hard to see that for any fixed $\alpha, \gamma \subseteq \{1, \dots, n\}$, $\sum_{\alpha \subseteq \beta \subseteq \gamma} (-1)^{|\beta| - |\alpha|} = \delta_{\alpha, \gamma}$. Assume that (3.1.1) holds. Then

$$\begin{aligned} \sum_{\alpha \subseteq \beta} (-1)^{|\beta| - |\alpha|} f(\beta) &= \sum_{\alpha \subseteq \beta} (-1)^{|\beta| - |\alpha|} \sum_{\beta \subseteq \gamma} g(\gamma) \\ &= \sum_{\gamma} g(\gamma) \sum_{\alpha \subseteq \beta \subseteq \gamma} (-1)^{|\beta| - |\alpha|} = g(\alpha). \end{aligned}$$

The other direction is proved similarly. □

The linear system (3.1.2), as the inverse to (3.1.1) is called *Inclusion-Exclusion*. It is most commonly used in the following enumerative context: Suppose that we have a finite set \mathcal{S} , and subsets A_1, \dots, A_n . For $\alpha \subseteq \{1, \dots, n\}$, let

$$f(\alpha) = \# \left(\bigcap_{i \in \alpha} A_i \right),$$

and

$$g(\alpha) = \# \left(\left(\bigcap_{i \in \alpha} A_i \right) \cap \left(\bigcap_{i \notin \alpha} \mathcal{S} \setminus A_i \right) \right).$$

We often describe A_i as being the set of elements in \mathcal{S} having a certain property P_i . Then $f(\alpha)$ is the number of elements in \mathcal{S} with properties P_i , for all $i \in \alpha$ (and possibly other properties), whereas $g(\alpha)$ is the number of elements in \mathcal{S} which have precisely the properties P_i , $i \in \alpha$, and no others. Clearly equation (3.1.1) holds, since the elements counted by $f(\alpha)$ must have properties precisely in some subset β where $\alpha \subseteq \beta$.

In this context, when $\alpha = \emptyset$, (3.1.2) states:

$$g(\emptyset) = \#(\mathcal{S} \setminus \bigcup_{i=1}^n A_i) = \sum_{\beta \subseteq [n]} (-1)^{|\beta|} \# \left(\bigcap_{i \in \beta} A_i \right) \quad (3.1.3)$$

which is equivalent to

$$\# \left(\bigcup_{i=1}^n A_i \right) = \sum_{\emptyset \neq \beta \subseteq [n]} (-1)^{|\beta|-1} \# \left(\bigcap_{i \in \beta} A_i \right).$$

Example 3.1.2. Let D_n be the number of permutations of $\{1, \dots, n\}$ with no fixed points (these are often called *derangements*). To count these, let $\mathcal{S} = S_n$ be the set of all permutations of $\{1, \dots, n\}$, and let subset A_i consisting of all the permutations that have i as a fixed point, for $i = 1, \dots, n$. Then, in the notation above, we have

$$D_n = g(\emptyset) = \sum_{\beta \subseteq \{1, \dots, n\}} (-1)^{|\beta|} f(\beta).$$

But, for any β of size k , we have $f(\beta) = (n - k)!$, and there are $\binom{n}{k}$ choices for such a β , which gives

$$D_n = \sum_{k=0}^n (-1)^k \binom{n}{k} (n - k)! = n! \sum_{k=0}^n \frac{(-1)^k}{k!}.$$

(This formula can also be obtained using the the theory of exponential generating functions.)

The derangement problem is a special case of permutations with *restricted positions*. We can think of a permutation σ as a collection of n non-attacking rooks on an $n \times n$ chessboard (we place a rook in row i , column $\sigma(i)$, for $i = 1, \dots, n$). If we let B be a set of “disallowed squares” on the board we can consider the set of all configurations of n rooks which avoid the squares in B . For example, in the derangement problem B is the set of diagonal squares, since we want permutations that avoid $\sigma(i) = i$.

Example 3.1.3. Let M_n be the number of permutations σ of $\{1, \dots, n\}$ for which which $\sigma(i) \neq i, i + 1 \pmod{n}$, for $i = 1, \dots, n$. This is called the *Ménage problem*. The historical context is as a problem in which n female-male couples are to be seated at a circular table so that no one is beside someone of their own gender, nor beside their mate. Here, the board of disallowed positions B has $2n$ squares, as in Figure 3.1.1. Let A_i , $i = 1, \dots, 2n$ be the




Figure 3.1.1: The disallowed board for the Ménage problem.

number of arrangements avoiding the i^{th} square B . Then

$$\begin{aligned}
 M_n &= \#(S_n \setminus \bigcup_{i=1}^n A_i) \\
 &= \sum_{\beta \subseteq \{1, \dots, 2n\}} (-1)^{|\beta|} \#(\bigcap_{i \in \beta} A_i) \\
 &= \sum_{k=0}^n (-1)^k (n-k)! R_k,
 \end{aligned}$$

where R_k is the number of ways to place k non-attacking rooks within the squares of B .

If we index the squares of B as in Figure 3.1.1, then R_k is precisely the number of k -subsets of $\{1, \dots, 2n\}$ with no pair of circularly consecutive elements. To enumerate these is an elementary counting exercise: $R_k = \frac{2n}{2n-k} \binom{2n-k}{k}$, and so

$$M_n = \sum_{k=0}^n (-1)^k \frac{2n}{2n-k} \binom{2n-k}{k} (n-k)!.$$

The first part of the argument above, works for any configuration of disallowed squares. For any B , the number of legal configurations of rooks is $\sum_{k=0}^n (-1)^k (n-k)! R_k$, where R_k is the number of ways to place k non-attacking rooks within the squares of B . However, in other examples, computing R_k may be difficult. An alternate approach (which can be equally computationally difficult) involves the MacMahon Master theorem (see Goulden & Jackson pp. 281–283).

3.2 Inclusion-Exclusion with generating functions

Now, suppose the set \mathcal{S} has a weight function $\text{wt} : \mathcal{S} \rightarrow \mathbb{Z}$. Let A_1, A_2, A_3, \dots be a finite or infinite list of subsets of \mathcal{S} . If the list is infinite, we assume that for any given n , there are only finitely many A_i that contain an element of weight n . Again, we often think of A_i as being the set of elements of \mathcal{S} that have a certain property P_i .

Instead of counting the number of elements in various intersections, we'll look at their generating functions with respect to the weight function wt . (Each subset of $\mathcal{T} \subset \mathcal{S}$ is also a set with weight function wt , and so we can consider its generating function $\sum_{\sigma \in \mathcal{T}} x^{\text{wt}(\sigma)}$.) For any finite subset $\alpha \subset \{1, 2, 3, \dots\}$, let $F_\alpha(x)$ be the generating function for the set $\bigcap_{i \in \alpha} A_i$. Let $G_\alpha(x)$ be the generating function for the subset $(\bigcap_{i \in \alpha} A_i) \cap (\bigcap_{i \notin \alpha} \mathcal{S} \setminus A_i)$. Let

$$F(x; u) = \sum_{\alpha} F_\alpha(x) u^{|\alpha|}$$

$$G(x; u) = \sum_{\alpha} G_\alpha(x) u^{|\alpha|}$$

Note that $G_\emptyset(x) = [u^0]G(x; u) = G(x; 0)$ is the generating function for the subset of \mathcal{S} with none of the properties P_1, P_2, \dots .

Theorem 3.2.1 (Generating function version of inclusion-exclusion). $G(x; u) = F(x; u - 1)$.

Note: the finiteness assumption we made on the subset A_i guarantees that $[x^n]F(x; u)$ is a polynomial in u for all n , so the substitution makes sense. We'll leave the proof and the generalization to multivariate generating functions as an exercise.

Example 3.2.2. Consider the set of all binary strings that do not have 011 as a substring. To compute the generating function for these (weighted by length), we let \mathcal{S} be the set of all binary strings, and let A_i be the set of all substrings of the form $* \cdots * 011 * \cdots *$, where the indicated 011 starts at the i^{th} position.

Then $[u^k]F(x; u)$ is the generating function for all binary strings with k marked occurrences of 011. For example,

$$\boxed{011}01011\boxed{011}1101101 \quad \text{and} \quad 01101\boxed{011}\boxed{011}1101101$$

are two different examples of binary strings with two marked occurrences of 011, so both of these contribute to $[u^2]F(x; u)$. The regular expression for binary strings with marked occurrences of 011 is $\{0, 1, \boxed{011}\}^*$, which gives us

$$F(x; u) = \frac{1}{1 - 2x - ux^3}.$$

Thus the generating function we want is

$$G_\emptyset(x) = G(x; 0) = F(x; -1) = (1 - 2x + x^3)^{-1}.$$

3.3 Sign reversing bijections

We now turn to one of the most important ideas for relating signed formulae and unsigned formulae in a combinatorial framework.

Let \mathcal{S} be a set of combinatorial objects. Let R be a ring (or sometimes just an abelian group) and suppose we have a “weight function” $\text{WT} : \mathcal{S} \rightarrow R$. Suppose we also have an “index function” $\text{ind} : \mathcal{S} \rightarrow \mathbb{Z}$. With this data, we can form a signed generating series for \mathcal{S} ,

$$S = \sum_{\sigma \in \mathcal{S}} (-1)^{\text{ind}(\sigma)} \text{WT}(\sigma).$$

Example 3.3.1. Let $\text{wt}_1, \dots, \text{wt}_k : \mathcal{S} \rightarrow \mathbb{Z}$ are (ordinary) weight functions on \mathcal{S} . A typical example of R and WT for this context will be $R = \mathbb{Q}[[x_1, \dots, x_k]]$, and $\text{WT}(\sigma) = x_1^{\text{wt}_1(\sigma)} \dots x_k^{\text{wt}_k(\sigma)}$. If $\text{ind}(\sigma) = 0$ for all $\sigma \in \mathcal{S}$, then then S is the ordinary multivariate generating function for \mathcal{S} .

Now suppose we have a subset $\mathcal{T} \subset \mathcal{S}$. We can consider the corresponding generating series for \mathcal{T} , $T = \sum_{\sigma \in \mathcal{T}} (-1)^{\text{ind}(\sigma)} \text{WT}(\sigma)$. Often, $\text{ind}(\sigma)$ is even for all $\sigma \in \mathcal{T}$, and so

$$T = \sum_{\sigma \in \mathcal{T}} \text{WT}(\sigma),$$

which does not have any signs.

Proposition 3.3.2. *Suppose there exists a bijection $\alpha : \mathcal{S} \setminus \mathcal{T} \rightarrow \mathcal{S} \setminus \mathcal{T}$ with the properties that $\text{WT}(\sigma) = \text{WT}(\alpha(\sigma))$, and $\text{ind}(\sigma)$ is even iff $\text{ind}(\alpha(\sigma))$ is odd. Then $S = T$.*

Proof. We have

$$\begin{aligned} \sum_{\sigma \in \mathcal{S} \setminus \mathcal{T}} (-1)^{\text{ind}(\sigma)} \text{WT}(\sigma) &= \sum_{\substack{\text{cycles} \\ \gamma \text{ of } \alpha}} \sum_{\sigma \in \gamma} (-1)^{\text{ind}(\sigma)} \text{WT}(\sigma) \\ &= \sum_{\gamma} \frac{1}{2} \left(\sum_{\sigma \in \gamma} (-1)^{\text{ind}(\sigma)} \text{WT}(\sigma) + \sum_{\sigma \in \gamma} (-1)^{\text{ind}(\alpha(\sigma))} \text{WT}(\alpha(\sigma)) \right) \\ &= 0, \end{aligned}$$

from which the result follows. □

The map α is called a *sign reversing bijection*. Often, we will have $\alpha = \alpha^{-1}$, in which case α is called a *sign reversing involution*. The tricky part, in any combinatorial setting, is to find a suitable α .

Example 3.3.3. As a first example, we give another proof of the inclusion-exclusion formula. Let \mathcal{S} , A_1, \dots, A_n , and $f(\alpha), g(\alpha)$ be as in Section 3.1. Let $\mathcal{S}' = \{(x, \alpha) \mid x \in A_i \text{ for all } i \in \alpha\}$. Define index and weight functions on \mathcal{S}' :

$$\text{ind}(x, \alpha) = |\alpha| \quad \text{WT}(x, \alpha) = 1.$$

Let $\mathcal{T} = \mathcal{S} \setminus (A_1 \cup \dots \cup A_n)$, and $\mathcal{T}' = \{(x, \alpha) \in \mathcal{S}' \mid x \in \mathcal{T}\}$. Then

$$\sum_{(x, \alpha) \in \mathcal{T}'} (-1)^{\text{ind}(x, \alpha)} \text{WT}(x, \alpha) = g(\emptyset)$$

and

$$\sum_{(x, \alpha) \in \mathcal{S}'} (-1)^{\text{ind}(x, \alpha)} \text{WT}(x, \alpha) = \sum_{\alpha \subseteq [n]} (-1)^{|\alpha|} f(\alpha).$$

To show that these are equal, we give a sign reversing involution σ on $\mathcal{S}' \setminus \mathcal{T}'$. For $(x, \alpha) \in \mathcal{S}' \setminus \mathcal{T}'$, we have $x \in A_i$ for some i . Let j be the smallest number such that $x \in A_j$. Let

$$\sigma(x, \alpha) = \begin{cases} (x, \alpha \setminus \{j\}) & \text{if } j \in \alpha \\ (x, \alpha \cup \{j\}) & \text{if } j \notin \alpha \end{cases}$$

This is weight preserving (everything has the same weight), sign-reversing (since index changes by ± 1), and an involution; hence $g(\emptyset) = \sum_{\alpha \subseteq [n]} (-1)^{|\alpha|} f(\alpha)$, as required.

For a second example, we prove the Gessel-Viennot theorem, which gives a formula for enumerating tuples of non-intersecting paths in a directed planar graph.

Theorem 3.3.4 (Gessel-Viennot). *Let G be a finite acyclic directed graph, embedded in the plane. Assign to each edge $e \in \text{edge}(G)$ a weight x_e in a commutative ring R . For a subgraph $H \subset G$, let $x^H = \prod_{e \in \text{edge}(H)} x_e$ be the product of the weights of the edges in H .*

Let $A_1, \dots, A_n, Z_n, \dots, Z_1$ be pairwise distinct vertices appearing (in that order) on the outer face of G . Define a matrix $M_{ij} = \sum_{P: A_i \rightarrow Z_j} x^P$ where the sum is taken over all (directed) paths from A_i to Z_j .

Define \mathcal{T} to be the set of subgraphs G consisting of a union of pairwise vertex-disjoint paths $P_1 \sqcup \dots \sqcup P_n$ where P_i is a path from A_i to Z_i . Then

$$\sum_{H \in \mathcal{T}} x^H = \det(M).$$

Example 3.3.5. In Figure 3.3.1 we have drawn a graph G with points $A_1, A_2, A_3, Z_3, Z_2, Z_1$ along the outer face. An example of a subgraph $H \in \mathcal{T}$, consisting of a vertex-disjoint union of non-intersecting paths, is highlighted. If we assign weight 1 to each edge of G , then M_{ij} is just the number of paths from A_i to Z_j , which is $\binom{6+i-j}{2+i-j}$. Thus the number of subgraphs in \mathcal{T} is given by

$$\det(M) = \det \begin{pmatrix} \binom{6}{2} & \binom{5}{1} & \binom{4}{0} \\ \binom{7}{3} & \binom{6}{2} & \binom{5}{1} \\ \binom{8}{4} & \binom{7}{3} & \binom{6}{2} \end{pmatrix} = 50,$$

which can be verified by listing all the possibilities.

Proof. Let \mathcal{S} be the set of all n -tuples of paths (P_1, \dots, P_n) where $P_i : A_i \rightarrow Z_{\pi(i)}$ for some permutation $\pi \in S_n$. Define

$$\begin{aligned} \text{WT}(P_1, \dots, P_n) &= x^{P_1} \dots x^{P_n} \\ \text{ind}(P_1, \dots, P_n) &= \ell(\pi) \end{aligned}$$




Figure 3.3.1: A directed acyclic graph, with a family of non-intersecting paths from $A_i \rightarrow Z_i$.

where $\ell(\pi)$ is the *length* of the permutation π . This is defined to be $\ell(\pi) = \#\{(i, j) \mid i < j \text{ and } \pi(i) > \pi(j)\}$. Then $\mathcal{T} \subset \mathcal{S}$ can be viewed as the subset of non-intersecting n -tuples of paths. We have

$$\begin{aligned} \sum_{(P_1, \dots, P_n) \subset \mathcal{S}} (-1)^{\text{ind}(P_1, \dots, P_n)} x^{P_1} \dots x^{P_n} &= \sum_{\pi \in S_n} (-1)^{\ell(\pi)} \prod_{i=1}^n \left(\sum_{P_i: A_i \rightarrow Z_{\pi(i)}} x^{P_i} \right) \\ &= \sum_{\pi \in S_n} (-1)^{\ell(\pi)} M_{1, \pi(1)} M_{2, \pi(2)} \dots M_{n, \pi(n)} \\ &= \det(M) \end{aligned}$$

Thus to prove the theorem, we construct a sign-reversing involution on $\mathcal{S} \setminus \mathcal{T}$, the non-disjoint n -tuples of paths (P_1, \dots, P_n) .

Given such a (P_1, \dots, P_n) , let i be the smallest index such that P_i meets another path. Let v be the first vertex on P_i , which is in another path. Let $j > i$ be the smallest index such that $v \in P_j$. Thus we can write $P_i = u_1 \dots v \dots u_s$, $P_j = w_1 \dots v \dots w_r$. We define

$$\alpha(P_1, \dots, P_n) = (Q_1, \dots, Q_n)$$

where $Q_k = P_k$ if $k \neq i, j$. $Q_i = u_1 \dots v \dots w_r$ is the path P_i up until v , and then P_j after v . Similarly $Q_j = w_1 \dots v \dots u_s$ is the path P_j up until v and then P_i after v .

Now we check that α has the desired properties. We have $\text{WT}(Q_1, \dots, Q_n) = \text{WT}(P_1, \dots, P_n)$, since exactly the same edges are used in the two tuples of paths. Moreover, $\text{ind}(P_1, \dots, P_n)$ is odd \iff $\text{ind}(Q_1, \dots, Q_n)$ is even, since the permutation associate to (P_1, \dots, P_n) gets multiplied by $(i \ j)$, which changes the length by an odd number. Finally, note that if we run the above construction on (Q_1, \dots, Q_n) , we obtain the same i, j and v , so $\alpha(Q_1, \dots, Q_n) = (P_1, \dots, P_n)$; hence α is an involution and the result follows from Proposition 3.3.2. \square

3.4 Möbius inversion

Inclusion-exclusion is actually a special case of a more general theorem on partially ordered sets.

A *partially ordered set* (or *poset*) \mathcal{P} is a set (also referred to as \mathcal{P}) with a binary relation \leq (or $\leq_{\mathcal{P}}$ if ambiguous), such that

- $x \leq x$ for all $x \in \mathcal{P}$. (Reflexive)
- If $x \leq y$ and $y \leq x$ then $y = x$. (Antisymmetric)
- If $x \leq y$ and $y \leq z$ then $x \leq z$. (Transitive)

We use $x < y$ to mean $x \leq y$ and $x \neq y$. We say that two elements x, y are *comparable* if $x \leq y$ or $y \leq x$; if not, then they are *incomparable*.

Example 3.4.1. Here are some examples of posets.

1. $\mathcal{P} = \{1, \dots, n\}$, with $i \leq_{\mathcal{P}} j$ when i is less than or equal to j . Here, every pair of elements is comparable, and this is called a *total order*.
2. $\mathcal{P} = \{1, \dots, n\}$, with $i \leq_{\mathcal{P}} j$ when $i|j$. We refer to this partial order as positive integers ordered by *divisibility*.
3. $\mathcal{P} = \{0, 1\}^{[n]}$. This can be viewed as the set of binary strings of length n with $\sigma_1 \dots \sigma_n \leq_{\mathcal{P}} \tau_1 \dots \tau_n$ if $\sigma_i \leq \tau_i$ for all i . We can also view this as the set of all subsets of $\{1, \dots, n\}$, with $\alpha \leq_{\mathcal{P}} \beta$ when $\alpha \subseteq \beta$. We refer to this partial order as subsets ordered by *inclusion*.
4. \mathcal{P} = the set of set partitions of $\{1, \dots, n\}$. Here, we have $P \leq_{\mathcal{P}} Q$ when every element of P is a subset of some element of Q . We refer to this partial order as set partitions ordered by *refinement*.

If $x < y$ and no z exists with $x < z < y$, then we say that y *covers* x (often denoted $x \prec y$). Because of transitivity and reflexivity, the cover relations are enough to specify a finite poset, and so we can represent a finite poset as a directed graph, with arrows $x \rightarrow y$ if $x \prec y$. Usually we omit the arrows when we draw the graph, by adopting the convention that all $x \prec y$ if y is above x . This is called the *Hasse diagram* of the poset. In Figure 3.4.1, we give the Hasse diagrams for the posets 1–4 from Example 3.4.1 above.

A *linear extension* of a poset \mathcal{P} with n elements is a listing (x_1, x_2, \dots, x_n) of the elements of \mathcal{P} that is consistent with the partial order, i.e. such that $x_i <_{\mathcal{P}} x_j$ only if $i < j$. For example, in Example 2 above with $n = 9$, some of the linear extensions are given by 135792468, 123456789, 124836957. A *minimal* element in a poset is one that is not greater than any other element (this is a source in the directed graph model). For example, the posets given in Figure 3.4.1 all have a unique minimal element. Similarly a *maximal* element is one that is not less than any other element. An increasing sequence of elements in a poset is called a *chain*. A *maximal chain* is a path in directed graph model from a minimal element to a maximal element.

For all $x, y \in \mathcal{P}$, define the *Möbius function* $\mu(x, y)$ recursively by




Figure 3.4.1: Examples of Hasse diagrams.

- $\mu(x, x) = 1$,
- $\mu(x, y) = - \sum_{x \leq z < y} \mu(x, z)$, if $x < y$,
- $\mu(x, y) = 0$ otherwise.

The fact that this uniquely defines μ easily follows by successively determining $\mu(x, y)$, for fixed x , with order of y determined by some linear extension of \mathcal{P} .

For example, in Figure 3.4.2, we have placed the value of $\mu(x, y)$ beside each vertex y , where x is the unique minimal element of the poset, for the posets given in Figure 3.4.1.




Figure 3.4.2: Examples of Möbius functions.

We also define the *poset incidence function* $\zeta(x, y)$, by

- $\zeta(x, y) = 1$, if $x \leq y$,
- $\zeta(x, y) = 0$ otherwise.

We can think of μ and ζ as square matrices, with rows and columns indexed by the elements of \mathcal{P} (in some order). The (x, y) -entry of μ is $\mu(x, y)$, and the (x, y) -entry of ζ is $\zeta(x, y)$. If we choose the order to be a linear extension of \mathcal{P} , then these matrices are both upper triangular.

Theorem 3.4.2. *Thinking of μ and ζ as matrices, $\mu = \zeta^{-1}$.*

Proof. We have

$$\begin{aligned} (\mu\zeta)_{x,y} &= \sum_{z \in \mathcal{P}} \mu(x, z)\zeta(z, y) \\ &= \sum_{x \leq z \leq y} \mu(x, z)\zeta(z, y) = \sum_{x \leq z \leq y} \mu(x, z), \end{aligned}$$

and the result follows from the definition of the Möbius function μ above. \square

Corollary 3.4.3.

$$\mu(x, y) = - \sum_{x < z \leq y} \mu(z, y), \quad \text{for } x < y.$$

Proof. Since μ and ζ are inverses, we also have $\zeta\mu = I$, and thus, $x < y$ we have

$$\begin{aligned} 0 &= (\zeta\mu)_{x,y} = \sum_{z \in \mathcal{P}} \zeta(x, z)\mu(z, y) \\ &= \sum_{x \leq z \leq y} \zeta(x, z)\mu(z, y) = \sum_{x \leq z \leq y} \mu(z, y), \end{aligned}$$

and the result follows. \square

The next result is called *Möbius inversion* (the Möbius function and Möbius inversion are both named after the 19th Century German mathematician August Ferdinand Möbius, as is the Möbius strip and the Möbius transformation in projective geometry).

Corollary 3.4.4. *Let f and g be functions of \mathcal{P} . Then,*

$$f(x) = \sum_{y \leq x} g(y), \quad x \in \mathcal{P},$$

if and only if

$$g(x) = \sum_{y \leq x} f(y)\mu(y, x), \quad x \in \mathcal{P}.$$

Proof. Viewing f and g be row vectors indexed by \mathcal{P} (using the same ordering of elements as we used to view ζ and μ as matrices), the first linear system is equivalent to $f = g\zeta$, and the second linear system is equivalent to $g = f\mu$, and the result follows since $\mu = \zeta^{-1}$. \square

The following, similar result is called *dual Möbius inversion*.

Corollary 3.4.5.

$$f(x) = \sum_{y \geq x} g(y), \quad x \in \mathcal{P},$$

if and only if

$$g(x) = \sum_{y \geq x} \mu(x, y)f(y), \quad x \in \mathcal{P}.$$

Proof. The first linear system is equivalent to $f^t = \zeta g^t$, and the second linear system is equivalent to $g^t = \mu f^t$. \square

We now describe the Möbius functions for the posets 1–3 from Example 3.4.1.

1. For the positive integers totally ordered, we have $\mu(i, j) = 1$ if $i = j$, $\mu(i, j) = -1$ if $i = j - 1$, and $\mu(i, j) = 0$ otherwise. In this case Corollary 3.4.4 simply says that

$$f(i) = g(1) + \cdots + g(i), \quad i \geq 1$$

if and only if

$$g(i) = f(i) - f(i - 1), \quad i \geq 1 \quad (\text{with } f(0) = g(0))$$

2. For the positive integers ordered by divisibility, we have $\mu(i, j) = (-1)^m$ if $\frac{j}{i}$ (where i divides j) has m , distinct, prime divisors, and $\mu(i, j) = 0$ otherwise. Corollary 3.4.4 says that

$$f(n) = \sum_{d|n} g(d) \iff g(n) = \sum_{d|n} \mu(d, n) f(d).$$

3. For $\{0, 1\}^{[n]}$, the poset of subsets of $\{1, \dots, n\}$, ordered by inclusion, we have $\mu(\alpha, \beta) = (-1)^{|\beta| - |\alpha|}$, if $\alpha \subseteq \beta$, and $\mu(\alpha, \beta) = 0$ otherwise. In this case, Corollary 3.4.5 says that

$$f(\alpha) = \sum_{\alpha \subseteq \beta} g(\beta), \quad \alpha \subseteq \{1, \dots, n\}$$

if and only if

$$g(\alpha) = \sum_{\alpha \subseteq \beta} (-1)^{|\beta| - |\alpha|} f(\beta), \quad \alpha \subseteq \{1, \dots, n\}.$$

This is precisely inclusion-exclusion.

Chapter 4

Partitions and tableaux


In this chapter we look at some of the enumerative and algorithmic combinatorics associated to partitions. In particular, we will consider three algorithms on Young tableaux, whose properties are so spectacular that they can't help but be useful. Indeed, they will play an important role in the theory of symmetric functions, in the next chapter.

4.1 Partitions

A *partition* $\lambda = (\lambda_1, \dots, \lambda_d)$ is a decreasing sequence of positive integers $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d > 0$. The *size* of λ is $|\lambda| = \lambda_1 + \lambda_2 + \dots + \lambda_d$. If $|\lambda| = n$ we also say that λ is a partition of n and write $\lambda \vdash n$. The numbers $\lambda_1, \dots, \lambda_d$ are called the *parts* of λ , and d is the *number of parts* of λ (also sometimes called the length of λ).

We adopt the convention that a $\lambda_k = 0$ for $k > d$. We may sometimes accidentally write extra zeros when writing out the parts of λ : for example $(3, 3, 2, 0, 0, 0, 0)$ is just another way of writing $(3, 3, 2)$ —the number of parts is 3, not 7.

We represent λ by its Ferrers diagram (also called Young diagram) which has λ_i boxes in row i . For example:



is the partition $(5, 4, 4, 1, 1, 1)$, which we may also write as 544111 or 54^21^3 .

To enumerate partitions, weighted by their size, we note that we can have any number of parts of any size. If we only allowed parts of size i , the generating function would be $(1 - x^i)^{-1}$; since a partition can be represented by listing of the number of parts of each size, the generating function is the product of all these:

$$\prod_{i=1}^{\infty} (1 - x^i)^{-1}.$$


If we add a second weight function, the number of parts, a similar argument shows that the two variable generating function is

$$\prod_{i=1}^{\infty} (1 - tx^i)^{-1}.$$

Exercise 4.1.1. A partition $\lambda = (\lambda_1, \dots, \lambda_d)$ is said to have distinct parts if $\lambda_1 > \lambda_2 > \dots > \lambda_d > 0$. Show that the generating function for the set of partitions with distinct parts, weighted by size and number of parts is

$$\prod_{i=1}^{\infty} (1 + tx^i).$$


The *conjugate* of a partition λ is the partition λ^t , whose diagram is obtained by reflecting along the main diagonal. For example, the conjugate of $\lambda = 544111$ (shown above) is $\lambda^t = 63331$:



The bijection $\lambda \mapsto \lambda^t$ on partitions shows that the generating function for partitions with at most k parts (weighted by size) equals the generating function for partitions with largest part at most k , which is

$$\prod_{i=1}^k (1 - x^i)^{-1}.$$

Using a similar type of relationship one can enumerate symmetric partitions — those for which $\lambda = \lambda^t$. In this case the generating function equals the generating function for partitions with distinct parts, all of which are odd. To see this, consider hooks of λ for the diagonal squares. (The hook of a box α consists of α itself, all boxes below α and all boxes to the right of α .)



Now consider the set $\mathcal{P}_{k,l}$ of partitions with at most k parts, and largest part at most l . To give the generating function, $P_{k,l}(x)$, we make the following definitions:

- For $n \in \mathbb{Z}_{\geq 0}$, let $n_q = 1 + q + \dots + q^{n-1}$.
- Let $n!_q = n_q(n-1)_q \dots 2_q 1_q$.
- Let $\binom{n}{k}_q = \frac{n!_q}{k!_q(n-k)!_q}$.

These are called q -analogues of n , $n!$, and $\binom{n}{k}$ respectively.

Theorem 4.1.2. *The generating function for $\mathcal{P}_{k,l}$ is $P_{k,l}(q) = \binom{k+l}{k}_q$.*

One possible approach to proving this is to note that

$$\begin{aligned} P_{k,l}(q) &= [t^0] \prod_{i=1}^k (1 - tq^i)^{-1} + \cdots + [t^l] \prod_{i=1}^k (1 - tq^i)^{-1} \\ &= [t^l] \frac{1}{1-t} \prod_{i=1}^k (1 - tq^i)^{-1}. \end{aligned}$$

One can then find a recurrence that is satisfied by both this expression and $\binom{k+l}{k}_q$.

Exercise 4.1.3. Show that $P_{k,l}(q) = q^k P_{k,l-1}(q) + P_{k-1,l}(q)$ and that

$$\binom{n}{k}_q = q^k \binom{n-1}{k}_q + \binom{n-1}{k-1}_q.$$

Hence, prove Theorem 4.1.2.

We'll give a different proof here.

Proof. This is a polynomial identity, so it suffices to prove it for infinitely many values of q . We'll prove this when q is a prime power. In this case, there exists a finite field \mathbb{F}_q with q elements. We'll count k -dimensional subspace of the vector space \mathbb{F}_q^{k+l} in two different ways.

1. To specify a subspace, we can give a $k \times (k+l)$ matrix A with linearly independent rows. It's easy to enumerate these. The first row cannot be the zero vector, so there are $q^{k+l} - 1$ possibilities. The second row cannot be linearly dependent with the first, so there are $q^{k+l} - q$ possibilities. In general, the i^{th} row cannot be in the $(i-1)$ -dimensional subspace spanned by the first $i-1$ rows, so there are $q^{k+l} - q^{i-1}$ possibilities. In all, there are

$$(q^{k+l} - 1)(q^{k+l} - q) \cdots (q^{k+l} - q^{k-1})$$

such matrices. But now A and A' have the same row space if there is an invertible $k \times k$ matrix B such that $A' = BA$. B is unique if it exists, so to count the number of subspaces, we divide by the number of invertible $k \times k$ matrices. By the same argument, now with $l = 0$, there are

$$(q^k - 1)(q^k - q) \cdots (q^k - q^{k-1})$$

possibilities for B . Hence the number of subspaces is

$$\frac{(q^{k+l} - 1)(q^{k+l} - q) \cdots (q^{k+l} - q^{k-1})}{(q^k - 1)(q^k - q) \cdots (q^k - q^{k-1})} = \binom{k+l}{k}_q.$$

2. To specify a subspace uniquely, we can give a $k \times (k+l)$ matrix in row reduced echelon form. In such a matrix, there are k pivot columns, with zeros to the left of, above and below

each pivot. The remaining entries are free to be any entry from \mathbb{F}_q . For example, if a 3×8 matrix has pivots in columns 1, 4, 6 then the matrix is of the form:


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

where the stars can be any entry from \mathbb{F}_q . After deleting the pivot columns in such a diagram, the positions of the stars will form a (reflected) Ferrers diagram in $\mathcal{P}_{k,l}$ (and every such Ferrers diagram corresponds to some selection of columns). Since there are q choices for each $*$, the number of subspaces is


$$\sum_{\lambda \in \mathcal{P}_{k,l}} q^{|\lambda|} = P_{k,l}(q). \quad \square$$

4.2 Young tableaux

Partitions form a poset, ordered by inclusion of diagrams. We have $\mu \subseteq \lambda$ if $\mu_i \leq \lambda_i$ for all i . If $\mu \subseteq \lambda$, then we can form the skew diagram λ/μ , which consists of the boxes of λ that are not in μ . (λ and μ are assumed to have the same upper left corner). For example:




A *semistandard Young tableau* (SSYT) is a filling of the boxes of a Ferrers diagram λ (or a skew diagram λ/μ) with positive integer entries, such that the entries weakly increase along rows (from left to right) and strictly increase down columns. The set of all SSYT of shape λ is denoted $\text{SSYT}(\lambda)$. Here are two examples:



The content of a SSYT T is the sequence (c_1, c_2, c_3, \dots) , where c_i is the number of times i occurs as an entry of T . In the example above on the right the content is $(2, 3, 5, 1, 0, 1, 1)$. As with partitions, we adopt the convention that (c_1, c_2, \dots, c_d) means the same thing as $(c_1, c_2, \dots, c_d, 0, 0, 0, \dots)$.

We'll also sometimes be interested in *standard Young tableaux* (SYT). A tableau $T \in \text{SSYT}(\lambda)$ (or skew shape λ/μ) is standard if its entries are $1, 2, \dots, |\lambda|$. The set of all SYT of shape λ is denoted $\text{SYT}(\lambda)$. Here are two examples:



In the next chapter, we'll study the generating functions for $\text{SSYT}(\lambda)$, where the vector valued weight function is the content. These generating functions are called *Schur functions*.

Example 4.2.1. We give a few preliminary examples of Schur functions. To keep things simple for now, we'll restrict ourselves to SSYT where the largest entry is at most 4. Placing this type of cap on the entries gives a Schur function in finitely many variables (in this case x_1, \dots, x_4) whereas without the cap we get a Schur function in infinitely many variables.

1. If $\lambda = 1$, then there are four SSYT :

$$\boxed{1} \quad \boxed{2} \quad \boxed{3} \quad \boxed{4}$$

Thus the generating function is $x_1 + x_2 + x_3 + x_4$.

2. If $\lambda = 2$, then there are ten SSYT :

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

Thus the generating function is

$$x_1^2 + x_2^2 + x_3^2 + x_4^2 + x_1x_2 + x_1x_3 + x_1x_4 + x_2x_3 + x_2x_4 + x_3x_4.$$

This is the sum of all monomials of degree 2, and is an example of a *complete* symmetric function.

3. If $\lambda = 111$, there are again four SSYT :

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

and so the generating function is $x_1x_2x_3 + x_1x_2x_4 + x_1x_3x_4 + x_2x_3x_4$. This quantity appears as $[z^3](z + x_1)(z + x_2)(z + x_3)(z + x_4)$, and is called an *elementary* symmetric function.

Notice that all of these examples are symmetric in the variables x_1, \dots, x_4 ; that is, if we interchange any two variables, the generating function stays the same. This will be true of all generating functions for SSYT . However, beyond this unproven observation, we're not in a position to say anything too deep just yet. We need to develop some algorithms for studying tableaux, and some of the general theory of symmetric functions before we can really gain any insight into Schur functions. We'll discuss three tableau algorithms in the rest of this chapter, and symmetric function theory in the next.

4.3 Row-Insertion

The first important algorithm we'll study is called *row-insertion*. This algorithm takes as input a tableau $T \in SSYT(\lambda)$ and a non-negative integer, and returns a tableau denoted $T \leftarrow a$. This new tableau has one extra box, and one extra a amongst its entries. We'll describe the algorithm as a process that modifies the tableau T .

The main step in the algorithm is a subroutine, which we'll call "row-merge". This takes as input a row of T , and a number b . If b is greater than or equal to every entry in the row, we add b to the end of the row, and STOP. Otherwise, let c be the leftmost entry of the row that is greater than b . Replace c with b and output c .

We produce $T \leftarrow a$ by first row-merging a into the first row of T . Then we row-merge the output into the second row of T , and so forth until we hit STOP. $T \leftarrow a$ is the result of this process.

Example 4.3.1. Let's row-insert 3 into the tableau shown below. We start by row-merging 3 into the first row.

1	1	3	4	4	9	← 3
2	4	4	5			
6						

The 3 is too small to go in the end of the first row, so it "bumps" the leftmost 4, which we row-merge into the second row.

1	1	3	3	4	9		
2	4	4	5				← 4
6							

The four bumps the 5

1	1	3	3	4	9		
2	4	4	4				
6						← 5	

which in turn bumps the 6, which finally settles on its own in the fourth row.

1	1	3	3	4	9	
2	4	4	4			
5						
6						

Proposition 4.3.2. $T \leftarrow a$ is a semistandard Young tableau.

Proof. We need to check that each time we bump, the newly inserted entry is (a) \geq the entry to its left; (b) \leq the entry to its right; (c) $<$ the entry below; (d) $>$ the entry above. Items (a), (b) and (c) are immediate from the construction. To see (d), note that the sequence of bumps moves weakly to the left at each stage. This implies that an entry gets inserted either below entries that were formerly to its left or below the previously inserted entry, which are smaller. □

Row-insertion is not a one-to-one operation: if $T \leftarrow a = T' \leftarrow a'$, it is not necessarily the case that $T = T'$ and $a = a'$. In Example 4.3.1, the same result could have been produced by inserting 9 into the tableau

1	1	3	3	4
2	4	4	4	
5				
6				

However, if in addition, we know the shape of the original tableau T , we can uniquely recover T and a from $T \leftarrow a$.

Proposition 4.3.3. *Given $U \in \text{SSYT}(\lambda^+)$ and a partition $\lambda \subseteq \lambda^+$ with one fewer box, there is a unique $T \in \text{SSYT}(\lambda)$, and a unique positive integer a such that $U = T \leftarrow a$.*

Proof. Run the row-insertion algorithm in reverse, starting with the entry in λ^+/λ , which was the last to be inserted. □

Example 4.3.4. If $T \in \text{SSYT}(311)$ and $T \leftarrow a$ is the tableau

1	1	4
2	4	
5		

then the last step in the row-insertion algorithm must have been to place the 4 in row two. This 4 must have been bumped by the second 1 in row one, and so we conclude that $a = 1$ and

$$T = \begin{array}{|c|c|c|} \hline 1 & 4 & 4 \\ \hline 2 & & \\ \hline 5 & & \\ \hline \end{array}$$

Proposition 4.3.5. *Suppose T is a SSYT of shape λ , and $T \leftarrow a$ has shape λ^+ and $T \leftarrow a \leftarrow b$ has shape λ^{++} .*

- (i) *If $a \leq b$ then the box of λ^{++}/λ^+ is strictly right of and weakly above the box of λ^+/λ .*
- (ii) *If $a > b$ then the box of λ^{++}/λ^+ is strictly below of and weakly left of the box of λ^+/λ .*

Exercise 4.3.6. Prove Proposition 4.3.5.

The *word* of a semistandard Young tableau T is the string $w(T) = w_1 w_2 \dots w_n$, obtained by reading the entries in each row from right to left, starting with the top row and moving downward. The order in which we read the entries is called the *reading order*. For example, if

$$T = \begin{array}{|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 5 \\ \hline 6 & 6 & 7 & 7 & \\ \hline 8 & & & & \\ \hline 9 & & & & \\ \hline \end{array}$$

then $w(T) = 54321776689$. The word is defined in the same way if T is a SSYT of skew shape.

Exercise 4.3.7. If $T \in \text{SSYT}(\lambda)$ and $w(T) = w_1 w_2 \dots w_n$, prove that

$$T = \epsilon \leftarrow w_n \leftarrow w_{n-1} \leftarrow \dots \leftarrow w_1,$$

where ϵ is the empty tableau.

Row-insertion can be used to define a product on semistandard Young tableaux of straight shape. If $U \in \text{SSYT}(\mu)$ and $T \in \text{SSYT}(\lambda)$, and $w(T) = w_1 w_2 \dots w_n$, define

$$U * T = U \leftarrow w_n \leftarrow w_{n-1} \leftarrow \dots \leftarrow w_1.$$

Here's a remarkable and not at all obvious fact:

Theorem 4.3.8. *The tableaux product is associative.*

We are not in a position to prove this completely right now, but Exercise 4.4.5 outlines an approach, modulo a theorem that will be proved in the next chapter. Further details can be found in Fulton's "Young tableaux".

4.4 Sliding

Our second algorithm is called *sliding*. The algorithm takes as input a skew tableau $T \in \text{SSYT}(\lambda/\mu)$ plus a marked corner of μ . It returns a tableau $T' \in \text{SSYT}(\lambda'/\mu')$ and a marked corner outside λ' . The partition μ' is just μ with the first marked corner deleted; whereas λ is λ' with second marked corner added.

The algorithm proceeds as follows. The marked corner swaps repeatedly places with the box below it, or the box to its right, until neither exists. If both are exist, it swaps with whichever of the two has the smaller entry. If both entries are equal, it swaps with the box below. The resulting tableau is T' .

Example 4.4.1. Let T be and the marked corner (denoted \times) of μ be as below.

				1	1
\times	1	1	2	2	
2	2	2			
3	3	4			

There is a 2 below the \times and a 1 to its right; the 1 is smaller, so it swaps with the \times .

				1	1
1	\times	1	2	2	
2	2	2			
3	3	4			

Again we have a 2 below the \times and a 1 to its right, so again the \times swaps with the 1.

				1	1
1	1	\times	2	2	
2	2	2			
3	3	4			

Now we have a 2 below \times and a 2 to its right; in the case of a tie, the \times swaps places with the box below it.

				1	1
1	1	2	2	2	
2	2	\times			
3	3	4			

At this point, there is no box to the right, so \times must swap places with the 4 below it.

				1	1
1	1	2	2	2	
2	2	4			
3	3	\times			

Now there is no box below or to the right of \times so we stop, having produced T' and a marked outside corner.

Proposition 4.4.2. *The T' produced by the sliding algorithm is a SSYT.*


Proof. It is easy to check that at every stage of the algorithm the entries are weakly increasing along rows and strictly increasing down columns. Thus at the end, when there are no “holes”, we have a SSYT. \square

Proposition 4.4.3. *Sliding is uniquely reversible. To reverse, the marked outside corner swaps places with the entry to its left or above it. If both exist, it swaps with whichever entry is larger; if both are equal, it swaps with the box above.*

Proof. If we were to try to swap the other way, the entries would not be weakly increasing along rows, and strictly increasing down columns, which is a position that can never occur in the sliding algorithm. \square

There is a deep relationship between sliding and row-insertion, as the next exercise shows:

Exercise 4.4.4. If $T \in \text{SSYT}(\lambda)$, and a is a positive integer, form the skew SSYT depicted below.




Now perform a sequence of slides (with each of the inside corners in turn) until the tableau has straight shape. Prove that the result is $T \leftarrow a$.

More generally, we can start with any skew SSYT, and perform a sequence of slides using the inside corner boxes, until we obtain a SSYT of straight shape. This process is called rectification. In the case of Exercise 4.4.4, there is only one way to do this; however, in general there will be many possible choices for which corner to use at every step. Another remarkable theorem states that the final result does not depend on any of these choices.

This is harder to prove than one might at first expect—each of the intermediate stages of this procedure do depend on the choices involved in a complicated way. We’ll prove this in the next chapter.

Exercise 4.4.5. Generalize Exercises 4.3.7 and 4.4.4, to show that the tableau product $U * T$ can be computed by rectifying the skew tableau



Deduce that the tableau product is associative from the fact that rectification is independent of choices.

4.5 Crystal operators

Our third algorithm defines several operators on SSYT, called *crystal operators*. The name comes from the theory of quantum groups. As input we take a SSYT T (possibly of skew shape) and a positive integer a . The algorithm will attempt to define three tableaux $E_a(T)$, $F_a(T)$, and $R_a(T)$ of the same shape as T . The first two will not be defined for all T, a ; when they are undefined, we will write $E_a(T) = \emptyset$ or $F_a(T) = \emptyset$. $R_a(T)$ is always defined. The operators E_a , F_a and R_a are called *raising*, *lowering* and *reflection* operators respectively.

The algorithm begins by crossing out every entry of T not equal to a or $a + 1$. (Crossing out should not be confused with deleting. By crossing out, we mean putting a temporary mark on each of these entries. At the end of the algorithm, these marks will be erased.) Now reading through remaining uncrossed entries in the reading order we proceed to cross out a ’s and $(a + 1)$ ’s in pairs as follows. When we read an $a + 1$ we look for the last read a , that has not already been crossed out. If it exists, we cross out the pair; otherwise we continue reading.

At the end, the sequence uncrossed entries will read

$$\underbrace{a + 1, \dots, a + 1}_s, \underbrace{a, \dots, a}_t$$

in the reading order for some $s, t \geq 0$. The operators E_a , F_a and R_a act by changing the number of a ’s and $(a + 1)$ ’s in this sequence. If $s \geq 1$, we obtain $E_a(T)$ by changing rightmost $a + 1$ in this sequence to a ; if $s = 0$ then $E_a(T) = \emptyset$. Similarly $F_a(T)$ is obtained by changing the leftmost a to $a + 1$, if $t \geq 1$, and $F_a(T) = \emptyset$ if $t = 0$. $R_a(T)$ is obtained by replacing the uncrossed entries (in sequence) by

$$\underbrace{a + 1, \dots, a + 1}_t, \underbrace{a, \dots, a}_s.$$

Example 4.5.1. Let's compute $F_2(T)$ where T is the tableau

				1	2	2	2	2	2	3
		1	1	2	3	4	4			
	2	2	3	3	4	5				
3	4	4								

Start by crossing out all entries except the 2's and 3's.

				1	2	2	2	2	2	3
		1	1	2	3	4	4			
	2	2	3	3	4	5				
3	4	4								

Now cross out 2's and 3's in pairs.

				1	2	2	2	2	2	3
		1	1	2	3	4	4			
	2	2	3	3	4	5				
3	4	4								

→

				1	2	2	2	2	2	3
		1	1	2	3	4	4			
	2	2	3	3	4	5				
3	4	4								

→

				1	2	2	2	2	2	3
		1	1	2	3	4	4			
	2	2	3	3	4	5				
3	4	4								

→

				1	2	2	2	2	2	3
		1	1	2	3	4	4			
	2	2	3	3	4	5				
3	4	4								

There are four 2's and one 3 remaining. The operator F_2 changes the first 2 to a 3. Then we erase the crossing-out marks. Hence $F_2(T)$ is the tableau on the right below.

				1	2	2	2	2	3	3
		1	1	2	3	4	4			
	2	2	3	3	4	5				
3	4	4								

→

				1	2	2	2	2	3	3
		1	1	2	3	4	4			
	2	2	3	3	4	5				
3	4	4								

Proposition 4.5.2. When they are defined, $E_a(T)$, $F_a(T)$, and $R_a(T)$ are SSYT.

Proof. If a and $a + 1$ are in the same column, then it is not hard to see that they both get crossed out; hence the columns are strictly increasing. Now note that the crossed out a 's are leftmost in their row, while the crossed out $(a + 1)$'s are rightmost in their row, from which we see that the rows are weakly increasing. □

Proposition 4.5.3. $R_a : \text{SSYT}(\lambda/\mu) \rightarrow \text{SSYT}(\lambda/\mu)$ is an involution.

Proof. One checks that the crossed out entries are the same for T and $R_a(T)$. □

This time, we'll actually prove a remarkable theorem about the algorithm.

Theorem 4.5.4. The crystal operators E_a , F_a , and R_a commute with sliding.

That is, if we apply E_a (or either of the others) and then perform a slide, the result is the same as performing the slide and then applying E_a .

Proof. Since E_a , F_a and R_a only affect boxes containing a and $a + 1$, it is enough to prove this for tableaux whose only entries are a and $a + 1$. For convenience, we'll assume $a = 1$, $a + 1 = 2$. We also only need to consider what happens to the entries that are right of and below the marked corner \times . There are three cases. In diagrams below, an entry that is shown to be crossed out, must be crossed out. There may be other crossed out entries too (including those which are shown as uncrossed), but these will not change the analysis.

Case 1: There is a 2 directly below \times , and there is a 1 to the right of \times that does not have a 2 below it.

$$\begin{array}{|c|c|c|c|c|c|c|c|} \hline \times & 1 & 1 & 1 & 1 & 2 & 2 & 2 \\ \hline 2 & 2 & 2 & & & & & \\ \hline \end{array} \rightarrow \begin{array}{|c|c|c|c|c|c|c|c|} \hline 1 & 1 & 1 & 1 & 1 & 2 & 2 & 2 \\ \hline 2 & 2 & 2 & & & & & \\ \hline \end{array}$$

In this case the \times moves to the right. Each 2 in the second row is paired up with a 1 in the first row, and the pair is crossed out. Changing any of the uncrossed entries does not affect the path of the \times , and sliding does not change which entries are crossed out. In other words, changing the uncrossed entries commutes with sliding.

Case 2: There is both a 1 and a 2 below \times .

$$\begin{array}{|c|c|c|c|c|c|c|c|} \hline \times & 1 & 1 & 1 & 1 & 2 & 2 & 2 \\ \hline 1 & 2 & 2 & 2 & & & & \\ \hline 2 & & & & & & & \\ \hline \end{array} \rightarrow \begin{array}{|c|c|c|c|c|c|c|c|} \hline 1 & 1 & 1 & 1 & 1 & 2 & 2 & 2 \\ \hline 2 & 2 & 2 & 2 & & & & \\ \hline \times & & & & & & & \\ \hline \end{array}$$

In this case, the \times moves downward and the analysis is the essentially same as the first case.

Case 3: If neither of the first two cases applies, we have either

$$\begin{array}{|c|c|c|c|c|c|c|c|} \hline \times & 1 & 1 & 1 & 2 & 2 & 2 & 2 \\ \hline 2 & 2 & 2 & 2 & & & & \\ \hline \end{array} \rightarrow \begin{array}{|c|c|c|c|c|c|c|c|} \hline 1 & 1 & 1 & 2 & 2 & 2 & 2 & 2 \\ \hline 2 & 2 & 2 & \times & & & & \\ \hline \end{array}$$

or

$$\begin{array}{|c|c|c|c|c|c|c|c|} \hline \times & 1 & 1 & 1 & 1 & 2 & 2 & 2 \\ \hline 1 & 2 & 2 & 2 & & & & \\ \hline \end{array} \rightarrow \begin{array}{|c|c|c|c|c|c|c|c|} \hline 1 & 1 & 1 & 1 & 1 & 2 & 2 & 2 \\ \hline 2 & 2 & 2 & \times & & & & \\ \hline \end{array}$$

In either case, the \times ends up the same place; thus again we see that changing any of the uncrossed entries commutes with sliding. \square

Chapter 5

Symmetric functions

In this chapter, we introduce the ring of symmetric functions, and discuss several of its interesting bases, including the Schur functions. We'll describe the change of basis formulas between the different bases, many of which have combinatorial significance. In the process, we'll see how the tableau algorithms from the last chapter can be used to uncover the deep structure of symmetric functions.

5.1 Symmetric functions in finitely many variables

Example 5.1.1. Here are some examples of symmetric functions in $\mathbb{Q}[x, y, z]$. The polynomials

$$e_1(x, y, z) = x + y + z, \quad e_2(x, y, z) = xy + xz + yx, \quad e_3(x, y, z) = xyz$$

are called the elementary symmetric functions in x, y, z . The polynomial

$$h_3(x, y, z) = x^3 + y^3 + z^3 + x^2y + y^2x + x^2z + zx^2 + y^2z + z^2y + xyz$$

is the complete (or homogeneous) symmetric function of degree 3. The polynomial

$$p_k(x, y, z) = x^k + y^k + z^k$$

is called a power sum symmetric polynomial of degree k . There are also many others, such as

$$3x^2 + 3y^2 + 3z^3 - x^3y^3z^3 + 14x^7y^7 + 14x^7z^7 + 14y^7z^7$$

that do not have special names. The term “symmetric” comes from the fact that if we interchange any two variables, we get the same function.

In general, the symmetric group S_d (permutations of $[d]$) acts on $\mathbb{Q}[x_1, \dots, x_d]$, by

$$(\sigma f)(x_1, \dots, x_d) = f(x_{\sigma(1)}, \dots, x_{\sigma(d)}).$$

We say that f is a symmetric function if $\sigma f = f$ for all $\sigma \in S_d$. We define

$$\Lambda^{(d)}(x) = \mathbb{Q}[x_1, \dots, x_d]^{S_d}$$

to be the set of all symmetric functions in $\mathbb{Q}[x_1, \dots, x_d]$. When the set of variables x_1, \dots, x_d is irrelevant, or is clear from context, we'll write $\Lambda^{(d)}$ for $\Lambda^{(d)}(x)$.

Since the sum and product of symmetric functions is symmetric, $\Lambda^{(d)}$ is a ring. It is also a vector space over \mathbb{Q} . As a vector space,

$$\Lambda^{(d)} = \bigoplus_{n \geq 0} \Lambda_n^{(d)}$$

where $\Lambda_n^{(d)}$ is the subspace of all homogeneous degree n symmetric polynomials. Each $\Lambda_n^{(d)}$ is a finite dimensional vector space. Our first task is to describe a basis for $\Lambda_n^{(d)}$.

Let $\lambda \vdash n$ be a partition with at most d parts. Define

$$m_\lambda = \sum x_{\sigma(1)}^{\lambda_1} x_{\sigma(2)}^{\lambda_2} \dots x_{\sigma(d)}^{\lambda_d}$$

where the sum is taken over all distinct monomials that arise if $\sigma \in S_d$ is a permutation. Also set $m_\lambda = 0$ if λ has more than d parts.

Example 5.1.2. If $d = 3$, then

$$\begin{aligned} m_{31} &= x_1^3 x_2 + x_1 x_3^3 + x_1^3 x_3 + x_1 x_2^3 + x_2^3 x_3 + x_2 x_3^3 \\ m_{222} &= x_1^2 x_2^2 x_3^2 \end{aligned}$$

Note that m_{222} has fewer than $3!$ terms; this is because each permutation σ produces the same monomial, and we do sum over repeated terms.

The symmetric functions m_λ are called *monomial symmetric functions*.

Proposition 5.1.3. *The set $\{m_\lambda \mid \lambda \vdash n, \lambda \text{ has at most } d \text{ parts}\}$ is a basis for $\Lambda^{(d)}$.*

Proof. The set $M_n = \{x_1^{\alpha_1} \dots x_d^{\alpha_d} \mid \alpha_1 + \dots + \alpha_d = n\}$ is a basis for $\mathbb{Q}[x_1, \dots, x_d]_n$. Since every monomial appears in exactly one m_λ , we see that $\{m_\lambda\}$ is linearly independent. To see that it is spanning, note that the map $P : f \mapsto \frac{1}{d!} \sum_{\sigma \in S_d} \sigma f$ is a projection from $\mathbb{Q}[x_1, \dots, x_d]_n$ onto $\Lambda_n^{(d)}$, and it sends each basis element of M_n to a scalar multiple of some m_λ . \square

Corollary 5.1.4. *$\dim \Lambda_n^{(d)}$ is the number of partitions of n with at most d parts.*

5.2 Symmetric functions in infinitely many variables

Notice that if $d \geq n$, then $\dim \Lambda_n^{(d)} = \#\{\lambda \vdash n\}$, which does not depend on d , so if d is large enough, the vector spaces $\Lambda_n^{(d)}, \Lambda_n^{(d+1)}, \Lambda_n^{(d+2)}, \dots$ are all isomorphic. In fact the ring structures are related too: For all d we have a ring homomorphism

$$\phi^{(d)} : \Lambda^{(d+1)} \rightarrow \Lambda^{(d)}, \quad f(x_1, \dots, x_d, x_{d+1}) \mapsto f(x_1, \dots, x_d, 0).$$

Under this homomorphism

$$\phi^{(d)}(m_\lambda) = m_\lambda$$

(where the m_λ on the left hand side is a function of $d + 1$ variables, and the m_λ on the right hand side is a function of d variables).

The upshot is that if we never look at polynomials of degree greater than d , then the rings $\Lambda^{(d)}$, $\Lambda^{(d+1)}$, $\Lambda^{(d+2)}$, \dots are isomorphic, for all practical purposes, i.e. if we consider the product

$$m_\mu m_\nu = \sum a_{\mu\nu}^\lambda m_\lambda \quad (5.2.1)$$

in $\Lambda^{(d)}$, then as long as $d \geq |\mu| + |\nu|$, the numbers $a_{\mu\nu}^\lambda$ do not depend on d .

We formalize this idea of taking the number of variables to be “large enough” by considering the ring of symmetric functions $\Lambda(x)$ in infinitely many variables x_1, x_2, \dots . Formally this is written as

$$\Lambda(x) = \varprojlim \Lambda^{(d)}(x).$$

As a vector space over \mathbb{Q} ,

$$\Lambda(x) = \bigoplus_{n \geq 0} \Lambda_n(x)$$

where $\Lambda_n(x)$ has a basis $\{m_\lambda \mid \lambda \vdash n\}$. We regard m_λ now as a formal sum

$$m_\lambda = \sum x_{\sigma(1)}^{\lambda_1} x_{\sigma(2)}^{\lambda_2} \cdots x_{\sigma(d)}^{\lambda_d}$$

taken over all distinct monomials that arise if $\sigma : [d] \rightarrow \mathbb{Z}_{>0}$ is an injective map. We can also add and multiply formally, but it is easier and more practical to note that the formula for multiplication is (5.2.1), the same as for a large enough finite number of variables.

Example 5.2.1. In infinitely many variables, we have

$$\begin{aligned} m_1 &= x_1 + x_2 + x_3 + \cdots \\ m_1^2 &= x_1^2 + 2x_1x_2 + x_2^2 + 2x_1x_3 + 2x_2x_3 + x_3^2 + \cdots \\ &= m_2 + 2m_{11}. \end{aligned}$$

Whereas with two variables we have

$$\begin{aligned} m_1 &= x_1 + x_2 \\ m_1^2 &= x_1^2 + 2x_1x_2 + x_2^2 \\ &= m_2 + 2m_{11} \end{aligned}$$

The formula is the same, since we never considered a function of degree greater than 2.

The monomial symmetric functions are one of the families of symmetric functions that will be important to us. Other important families are the elementary symmetric functions, the complete symmetric functions, the power sum symmetric functions, and the Schur functions. These are defined in both Λ and $\Lambda^{(d)}$.

The *elementary* symmetric functions are defined to be

$$e_n = m_{1^n} = \sum_{i_1 < i_2 < \cdots < i_n} x_{i_1} \cdots x_{i_n}.$$

The *complete* (or *homogeneous*) symmetric functions are

$$h_n = \sum_{\lambda \vdash n} m_\lambda = \sum_{i_1 \leq i_2 \leq \dots \leq i_n} x_{i_1} \cdots x_{i_n}.$$

The *power sum* symmetric functions are

$$p_n = m_n = \sum_i x_i^n.$$

If $\lambda = (\lambda_1, \dots, \lambda_d)$ is a partition with d parts, we put

$$e_\lambda = e_{\lambda_1} \cdots e_{\lambda_n}, \quad h_\lambda = h_{\lambda_1} \cdots h_{\lambda_n}, \quad p_\lambda = p_{\lambda_1} \cdots p_{\lambda_n}.$$

From time to time, we may also use these definitions when $\lambda = (\lambda_1, \dots, \lambda_d)$ is not a partition.

Exercise 5.2.2. Let $\lambda = (\lambda_1, \dots, \lambda_n)$ be a partition with d parts. Let \mathcal{P} be the set of all functions $f : [d] \rightarrow \mathbb{Z}_{>0}$. Define a weight function $\text{wt}^\lambda(f) = (\text{wt}_1^\lambda(f), \text{wt}_2^\lambda(f), \dots)$ on \mathcal{P} by

$$\text{wt}_i^\lambda(f) = \sum_{j:f(j)=i} \lambda_j.$$

Prove that the generating function for \mathcal{P} with respect to wt^λ is $p_\lambda(x_1, x_2, \dots)$.

The *Schur function* $s_\lambda \in \Lambda$ is defined to be the generating function for SSYT weighted by content. Hence

$$s_\lambda = \sum_{T \in \text{SSYT}(\lambda)} x_1^{c_1(T)} x_2^{c_2(T)} \cdots$$

where $(c_1(T), c_2(T), \dots)$ is the content of T . Right now, this is a formal infinite sum, but we'll see shortly that it is a symmetric function of degree $n = |\lambda|$, and hence is an element of the finite dimensional vector space Λ_n .

In $\Lambda^{(d)}$, the Schur function $s_\lambda(x_1, \dots, x_d)$ is the generating function for SSYT of shape λ with entries $\leq d$. This is also obtained by evaluating $s_\lambda(x_1, x_2, \dots)$ at $x_{d+1} = x_{d+2} = x_{d+3} = \dots = 0$.

Example 5.2.3. We've already seen some examples of Schur functions in $\Lambda^{(4)}$ (Example 4.2.1). For the same partitions, the Schur functions in Λ look similar:

- $s_1 = x_1 + x_2 + x_3 + \dots$
- $s_2 = x_1^2 + x_1x_2 + x_2^2 + x_1x_3 + x_2x_3 + x_3^2 + \dots = h_2$.
- $s_{111} = x_1x_2x_3 + x_1x_2x_4 + x_1x_3x_4 + x_2x_3x_4 + x_1x_2x_5 + \dots = e_3$.
- $s_{21} = x_1^2x_2 + x_1x_2^2 + 2x_1x_2x_3 + \dots = m_{21} + 2m_{111}$.

Exercise 5.2.4. Check that $s_n = h_n$ and $s_{1^n} = e_n$ for all n .

Proposition 5.2.5. The Schur function s_λ is a symmetric function.

Proof. It is enough to show, for each $i \geq 1$ that the number of SSYT of shape λ and content $(c_1, c_2, \dots, c_i, c_{i+1}, \dots)$ equals the number with content $(c_1, c_2, \dots, c_{i+1}, c_i, \dots)$. Since the crystal reflection operator R_i gives a bijection between these, the result follows immediately. \square

5.3 Five bases

Theorem 5.3.1. *The following are all bases for Λ_n : $\{s_\lambda \mid \lambda \vdash n\}$, $\{e_\lambda \mid \lambda \vdash n\}$, $\{h_\lambda \mid \lambda \vdash n\}$, $\{p_\lambda \mid \lambda \vdash n\}$.*

Hence we have five different bases for Λ : $\{m_\lambda\}$, $\{e_\lambda\}$, $\{h_\lambda\}$, $\{p_\lambda\}$ $\{s_\lambda\}$, where λ runs over all partitions.

This proof uses the concept of lexicographic order on sequences of integers. If $\alpha = (\alpha_1, \alpha_2, \dots)$ and $\beta = (\beta_1, \beta_2, \dots)$, we say that α is lexicographically before β , and write $\alpha <_{\text{lex}} \beta$ if $\alpha_1 < \beta_1$ or $\alpha_1 = \beta_1, \dots, \alpha_k = \beta_k$ and $\alpha_{k+1} < \beta_{k+1}$ for some k . This defines a total order on sequences. For example $(4, 4, 3, 1, 5, \dots) <_{\text{lex}} (4, 4, 3, 2, 0, 1, \dots)$.

Proof. We can write

$$s_\lambda = \sum_{\mu \vdash n} K_{\lambda\mu} m_\mu \quad (5.3.1)$$

From the definition of s_λ we have $K_{\lambda\mu} = [x_1^{\mu_1} \cdots x_d^{\mu_d}] s_\lambda$ is the number of SSYT of shape λ and content μ . We claim that $K_{\lambda\mu} = 0$ if $\lambda <_{\text{lex}} \mu$ and $K_{\lambda\lambda} = 1$.

To see this, note that if T is a SSYT of shape λ then every occurrence of entry i must be in row i or above. Thus $\mu_1 + \dots + \mu_i \leq \lambda_1 + \dots + \lambda_i$ for all i , with equality if and only if every entry in row j equals j for $j = 1, \dots, i$. This proves the claim.

But now see that (5.3.1) is an upper triangular system of linear equations,

$$s_\lambda = m_\lambda + \sum_{\mu <_{\text{lex}} \lambda} K_{\lambda\mu} m_\mu.$$

Hence, this is an invertible system, and it follows that $\{s_\lambda \mid \lambda \vdash n\}$ is a basis for Λ_n .

A similar type of argument works for $\{e_\lambda\}$ and $\{p_\lambda\}$. For $\{e_\lambda\}$ the upper triangular system is of the form

$$e_{\lambda^t} = m_\lambda + \sum_{\mu <_{\text{lex}} \lambda} L_{\lambda\mu} m_\mu.$$

To see this, we observe that if $\mu_1 > \lambda_1$ (λ_1 is the number of parts of λ^t) then $[x_1^{\mu_1}] e_{\lambda^t} = 0$; we proceed by induction if $\mu_1 = \lambda_1^t$.

For $\{p_\lambda\}$ the upper triangular system is

$$p_\lambda = \sum_{\mu \geq_{\text{lex}} \lambda} M_{\lambda\mu} m_\mu,$$

with $M_{\lambda\lambda} > 0$. To see this, use the interpretation of Exercise 5.2.2, and show that if $(\text{wt}_1^\lambda(f), \text{wt}_2^\lambda(f), \dots)$ is a partition, it must be lexicographically greater than or equal to λ .

For $\{h_\lambda\}$, we need a different argument. A linear dependency among $\{e_\lambda\}$ is the same thing as an algebraic relation among $\{e_1, e_2, \dots\}$; hence the latter is an algebraically independent set of generators in Λ . As such, we can uniquely define a homomorphism $\omega : \Lambda \rightarrow \Lambda$ by $\omega(e_i) = h_i$. In the ring $\Lambda[[t]]$, let

$$\begin{aligned} H(t) &= \sum_{k \geq 0} h_k t^k = \prod_{j \geq 1} (1 - x_j t)^{-1} \\ E(t) &= \sum_{k \geq 0} e_k t^k = \prod_{j \geq 1} (1 + x_j t) \end{aligned} \quad (5.3.2)$$

From the right hand sides, we see that

$$E(t)H(-t) = 1.$$

Applying ω to both sides ($\omega(t) = t$) we obtain $H(t)\omega(H(-t)) = 1$; replacing $t \rightarrow -t$ gives

$$H(-t)\omega(H(t)) = 1 = E(t)H(-t).$$

Hence $\omega(H(t)) = E(t)$ and $\omega(h_i) = e_i$. Thus we see that $\omega = \omega^{-1}$, which tells us that ω is an isomorphism, and therefore that $\{h_\lambda \mid \lambda \vdash n\}$ is also a basis for Λ_n . \square

Exercise 5.3.2. Fill in missing the details in the proofs that $\{e_\lambda\}$ and $\{p_\lambda\}$ are bases for Λ .

The involution $\omega : \Lambda \rightarrow \Lambda$ in this last proof is called the *fundamental involution*.

Proposition 5.3.3. *Applying the fundamental involution to p_k gives $\omega(p_k) = (-1)^{k-1}p_k$; hence $\omega(p_\lambda) = (-1)^{|\lambda|-\ell(\lambda)}p_\lambda$, where $\ell(\lambda)$ is the number of parts in λ .*

Proof. From (5.3.2), we have

$$\begin{aligned} \log H(t) &= \sum_{j \geq 1} \sum_{k \geq 1} x_j^k \frac{t^k}{k} = \sum_{k \geq 1} p_k \frac{t^k}{k} \\ \log E(t) &= \sum_{j \geq 1} \sum_{k \geq 1} x_j^k \frac{(-1)^{k-1} t^k}{k} = \sum_{k \geq 1} (-1)^{k-1} p_k \frac{t^k}{k}. \end{aligned} \tag{5.3.3}$$

Since $\omega(H(t)) = E(t)$, and ω is an isomorphism, we have $\omega(p_k) = (-1)^{k-1}p_k$. The second statement follows from the fact that $|\lambda| - \ell(\lambda)$ is congruent to the number of parts of even length in λ modulo 2. \square

Equations (5.3.2), and (5.3.3) can be used to explicitly write the symmetric functions e_i , h_i , p_i in terms of one another. For example for $n \geq 1$, we have $[t^n]E(t)H(-t) = [t^n]1 = 0$; hence

$$\sum_{k=0}^n (-1)^k e_k h_{n-k} = 0.$$

We can rewrite this as $h_n = \sum_{k=1}^n (-1)^{k-1} e_k h_{n-k}$, which by induction on n , allows us to write h_n in terms of e_1, \dots, e_n . Applying ω , this also gives e_n in terms of h_1, \dots, h_n .

Similarly, differentiating (5.3.3) with respect to t gives

$$\begin{aligned} \frac{H'(t)}{H(t)} &= \sum_{k \geq 1} p_k t^{k-1} \\ \sum_{n \geq 1} n h_n t^{n-1} &= \sum_{\ell \geq 0} h_\ell t^\ell \sum_{k \geq 1} p_k t^{k-1} \\ n h_n &= \sum_{k=1}^n h_{n-k} p_k, \quad n \geq 1, \end{aligned}$$

which allows us to write h_n in terms of p_1, \dots, p_n or vice-versa.

The numbers $K_{\lambda\mu}$ from (5.3.1) are called *Kostka numbers*. Our next goal is to show that the Kostka numbers also form a change of basis matrix relating h_λ to s_λ : specifically,

$$h_\mu = \sum_{\lambda \vdash n} K_{\lambda\mu} s_\lambda.$$

This relationship comes from type of duality on Λ , which we'll study next.

5.4 The Hall inner product

Define a bilinear form $\langle \cdot, \cdot \rangle$ on Λ (or equivalently on each Λ_n) by declaring

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

Thus if $f = \sum_\lambda c_\lambda h_\lambda$ and $g = \sum_\mu d_\mu m_\mu$ ($c_\lambda, d_\mu \in \mathbb{Q}$) are symmetric functions, then

$$\langle f, g \rangle = \sum_\lambda \sum_\mu c_\lambda d_\mu \langle h_\lambda, m_\mu \rangle = \sum_\lambda c_\lambda d_\lambda$$

Theorem 5.4.1. *Suppose that $\{u_\lambda\}$ and $\{v_\lambda\}$ are bases for Λ . Then the following are equivalent:*

(i) $\langle u_\lambda, v_\mu \rangle = \delta_{\lambda,\mu}$

(ii) $\sum_\lambda u_\lambda(x) v_\lambda(y) = \prod_{i,j \geq 1} (1 - x_i y_j)^{-1}$

Here $u_\lambda(x) = u_\lambda(x_1, x_2, \dots)$ and $v_\mu(y) = v_\mu(y_1, y_2, \dots)$.

Proof. First we note that

$$\prod_{i,j \geq 1} (1 - x_i y_j)^{-1} = \prod_{i,j \geq 1} \sum_{k_j \geq 0} h_{k_j}(x) y_j^{k_j} = \sum_\lambda m_\lambda(x) h_\lambda(y)$$

The rest of the proof is now just fiddling with linear algebra. Let $C_{u,v}(x, y) = \sum_\lambda u_\lambda(x) v_\lambda(y)$. Since $\langle \cdot, \cdot \rangle$ is non-degenerate, we have

$$C_{u,v}(x, y) = C_{h,m}(x, y) = \prod_{i,j \geq 1} (1 - x_i y_j)^{-1}$$

if and only if

$$\langle C_{u,v}(x, y), v_\mu(x) \rangle_x = \langle C_{h,m}(x, y), v_\mu(x) \rangle_x \quad \text{for all } \mu. \quad (5.4.1)$$

(By $\langle \cdot, \cdot \rangle_x$ we mean evaluate the bilinear form with respect to the variables x_1, x_2, \dots , treating y_1, y_2, \dots as scalars.)

We can evaluate the left hand side of (5.4.1) by observing that

$$\langle C_{h,m}(x, y), m_\mu(x) \rangle_x = \sum_\lambda m_\mu(y) \langle h_\lambda(x), m_\lambda(x) \rangle_x = m_\mu(y)$$

for all μ . Since $\{m_\mu(x)\}$ is a basis for $\Lambda(x)$ this implies that $\langle C_{h,m}(x,y), f(x) \rangle_x = f(y)$ for all $f \in \Lambda$. In particular

$$\langle C_{h,m}(x,y), v_\mu(x) \rangle_x = v_\mu(y).$$

But now the right hand side of (5.4.1) is

$$\langle C_{u,v}(x,y), v_\mu(x) \rangle_x = \sum_{\lambda} v_\lambda(y) \langle u_\lambda(x), v_\mu(x) \rangle_x = \sum_{\lambda} v_\lambda(y) \langle u_\lambda, v_\mu \rangle$$

which equals $v_\mu(y)$ if and only if $\langle u_\lambda, v_\mu \rangle = \delta_{\lambda,\mu}$. \square

A quantity that appears frequently when dealing with power sum symmetric functions is

$$z(\lambda) = 1^{i_1} 2^{i_2} \dots i_1! i_2! \dots,$$

where the partition λ has i_j parts equal to j , $j \geq 1$.

Exercise 5.4.2. If $\lambda \vdash n$, show that $\frac{n!}{z(\lambda)}$ is the number of permutations of $[n]$ whose cycles (in some order) have lengths $\lambda_1, \lambda_2, \dots$

Corollary 5.4.3. $\{p_\lambda\}$ is an orthogonal basis, with $\langle p_\lambda, p_\mu \rangle = \delta_{\lambda,\mu} z(\lambda)$.

Proof.

$$\begin{aligned} \prod_{i,j \geq 1} (1 - x_i y_j)^{-1} &= \exp \sum_{i,j \geq 1} \sum_{k \geq 1} \frac{x_i^k y_j^k}{k} \\ &= \exp \sum_{k \geq 1} \frac{p_k(x) p_k(y)}{k} = \prod_{k \geq 1} \exp \frac{p_k(x) p_k(y)}{k} \\ &= \prod_{k \geq 1} \sum_{i_k \geq 0} \frac{p_k(x)^{i_k} p_k(y)^{i_k}}{k^{i_k} i_k!} \\ &= \sum_{i_1, i_2, \dots \geq 0} \prod_{k \geq 1} \frac{p_k(x)^{i_k} p_k(y)^{i_k}}{k^{i_k} i_k!} = \sum_{\lambda} \frac{p_\lambda(x) p_\lambda(y)}{z(\lambda)}, \end{aligned}$$

from which we deduce $\langle \frac{p_\lambda}{z(\lambda)}, p_\mu \rangle = \delta_{\lambda,\mu}$. \square

Corollary 5.4.4. The bilinear form $\langle \cdot, \cdot \rangle$ is symmetric and positive definite. Moreover ω is an isometry, i.e. $\langle \omega(f), \omega(g) \rangle = \langle f, g \rangle$ for all $f, g \in \Lambda$

Proof. To prove this, it suffices to check these statements for the basis $\{p_\lambda\}$. Indeed, we have $\langle p_\lambda, p_\mu \rangle = \langle p_\mu, p_\lambda \rangle$, $\langle p_\lambda, p_\lambda \rangle > 0$, and, $\langle \omega(p_\lambda), \omega(p_\mu) \rangle = \langle \pm p_\lambda, \pm p_\mu \rangle = \langle p_\lambda, p_\mu \rangle$. \square

Hence $\langle \cdot, \cdot \rangle$ deserves to be called an inner product; it is known as the *Hall inner product*. The expression $\prod_{i,j \geq 1} (1 - x_i y_j)^{-1}$ is known as the *Cauchy Kernel*.

5.5 The Robinson-Schensted-Knuth correspondence

Theorem 5.5.1. *The basis of Schur functions $\{s_\lambda\}$ is an orthonormal basis for Λ .*

Proof. We'll show that

$$\sum_{\lambda} s_{\lambda}(x)s_{\lambda}(y) = \prod_{i,j \geq 1} (1 - x_i y_j)^{-1}.$$

The left hand side of this expression is the generating function for pairs of SSYT of the same shape, weighted by their contents:

$$s_{\lambda}(x)s_{\lambda}(y) = \sum_{P,Q \in \text{SSYT}(\lambda)} x^{\text{content}(Q)} y^{\text{content}(P)}$$

The right hand side is the generating function for multisets of ordered pairs of positive integers $\{(a_1, b_1), (a_2, b_2), \dots, (a_k, b_k)\}$, $k \geq 0$, where the weight is again “content”, i.e. $(c_1, c_2, \dots; d_1, d_2, \dots)$, where c_j is the number of i 's such that $a_i = j$, and d_j is the number of i 's such that $b_i = j$. Indeed we have

$$\begin{aligned} \prod_{i,j \geq 1} (1 - x_i y_j)^{-1} &= \sum_{\{(a_i, b_i)\}} x_{a_1} \cdots x_{a_k} y_{b_1} \cdots y_{b_k} \\ &= \sum_{\{(a_i, b_i)\}} x^{\text{content}(\{a_i\})} y^{\text{content}(\{b_i\})} \end{aligned}$$

To prove the theorem, we need to give a bijection between pairs (P, Q) of SSYT of the same shape, and multisets $\{(a_1, b_1), (a_2, b_2), \dots, (a_k, b_k)\}$, where a_1, \dots, a_k are the entries of Q , and b_1, \dots, b_k are the entries of P .

We do this using an algorithm based on row-insertion, called the *Robinson-Schensted-Knuth* (RSK) *correspondence*.

Assume that $(a_1, b_1), (a_2, b_2), \dots, (a_k, b_k)$ are lexicographically ordered; that is $a_i < a_{i+1}$, or $a_i = a_{i+1}$ and $b_i \leq b_{i+1}$. We create a sequence of SSYT P_i and Q_i , as follows. We begin with $P_0 = Q_0 = \epsilon$ to be the empty tableau. Then define $P_i = P_{i-1} \leftarrow b_i$. We obtain Q_i by adding a box to Q_{i-1} so that it has the same shape as P_i , and putting a_i in that box. The corresponding pair of tableaux are $P = P_k, Q = Q_k$. Note that Q has entries a_1, \dots, a_n , and P has entries b_1, \dots, b_n . (See Figure 5.5.1 for an example.)

First, we show that P and Q are both tableaux. For P , this follows from the fact that row-insertion creates a tableau. For Q , since the entries a_i are added in increasing order, the rows and columns are weakly increasing. To see that the columns are strictly increasing, note that if $a_i = a_{i+1} = \dots = a_j$, then $b_i \leq b_{i+1} \leq \dots \leq b_j$; by Proposition 4.3.5, the sequence of added boxes always moves strictly to the right whenever $a_i = a_{i+1} = \dots = a_j$.

Next, we show that the algorithm is reversible, which shows that the RSK correspondence is injective. To reverse it, we look for the largest entry in $Q = Q_k$. This is a_k . If there is more than one such entry, the rightmost must be the last to be added, which tells us Q_{k-1} is obtained by obtaining deleting this rightmost largest entry. This gives us the shape of P_{k-1} form which we reverse the last row insertion and determine b_k and P_{k-1} . By induction this determines all (a_i, b_i)

$P_1 = \boxed{1}$	$Q_1 = \boxed{1}$
$P_2 = \boxed{1\ 1}$	$Q_2 = \boxed{1\ 1}$
$P_3 = \boxed{1\ 1\ 3}$	$Q_3 = \boxed{1\ 1\ 1}$
$P_4 = \begin{array}{ c c c } \hline 1 & 1 & 2 \\ \hline 3 & & \\ \hline \end{array}$	$Q_4 = \begin{array}{ c c c } \hline 1 & 1 & 1 \\ \hline 2 & & \\ \hline \end{array}$
$P_5 = \begin{array}{ c c c c } \hline 1 & 1 & 2 & 6 \\ \hline 3 & & & \\ \hline \end{array}$	$Q_5 = \begin{array}{ c c c c } \hline 1 & 1 & 1 & 2 \\ \hline 2 & & & \\ \hline \end{array}$
$P_6 = \begin{array}{ c c c c } \hline 1 & 1 & 1 & 6 \\ \hline 2 & & & \\ \hline 3 & & & \\ \hline \end{array}$	$Q_6 = \begin{array}{ c c c c } \hline 1 & 1 & 1 & 2 \\ \hline 2 & & & \\ \hline 5 & & & \\ \hline \end{array}$
$P_7 = \begin{array}{ c c c c } \hline 1 & 1 & 1 & 2 \\ \hline 2 & 6 & & \\ \hline 3 & & & \\ \hline \end{array}$	$Q_7 = \begin{array}{ c c c c } \hline 1 & 1 & 1 & 2 \\ \hline 2 & 5 & & \\ \hline 5 & & & \\ \hline \end{array}$

Figure 5.5.1: The RSK correspondence for $\{(1, 1), (1, 1), (1, 3), (2, 2), (2, 6), 5, 1), (5, 2)\}$.

Finally, we show that every pair of SSYT of the same shape arises, which shows that the correspondence is surjective. To see this, just apply the reverse RSK algorithm to any pair (P, Q) and note that applying RSK to the resulting multiset of pairs $\{(a_1, b_1), \dots, (a_k, b_k)\}$ produces (P, Q) . \square

A special case of the RSK algorithm is the original Robinson-Schensted correspondence, in which the input is of the form $\{(1, \pi(1)), \dots, (n, \pi(n))\}$, where $\pi \in S_n$ is permutation of $[n]$. In this case we get a bijection between S_n and pairs of standard Young tableaux of the same shape. Letting $f^\lambda = \#\text{SYT}(\lambda)$, we obtain the identity

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

This identity also has an interpretation in representation theory of the symmetric group.

An interesting, non-obvious feature of the RSK correspondence is that if $\{(a_1, b_1), \dots, (a_n, b_n)\}$ corresponds to (P, Q) , then $\{(b_1, a_1), \dots, (b_n, a_n)\}$ corresponds to (Q, P) . We will not prove this. We will, however, note that in the case of the original Robinson-Schensted correspondence, this tells us that $\pi \in S_n$ corresponds to (P, Q) iff π^{-1} corresponds to (Q, P) . It follows that $P = Q$ if and only if $\pi = \pi^{-1}$ is an involution. By counting pairs (P, P) of equal standard Young tableaux, we obtain

$$\sum_{\lambda \vdash n} f^\lambda = \#\{\pi \in S_n \mid \pi = \pi^{-1}\}.$$

Involutions are easy to count: they are characterized by the fact that their cycles have

lengths 1 or 2. Hence the species of involutions is $\mathcal{E}[\mathcal{E}_1 \oplus \mathcal{E}_2]$, and

$$\sum_{\lambda \vdash n} f^\lambda = [x^n] e^{x + \frac{x^2}{2}}.$$

As a corollary of Theorem 5.5.1, we obtain one relationship between the Schur functions s_λ and the complete symmetric functions.

Corollary 5.5.2.

$$h_\mu = \sum_{\lambda \vdash n} K_{\lambda\mu} s_\lambda$$

Proof. Apply the operator $\sum_{\lambda} \langle \cdot, h_\mu \rangle s_\lambda$ to both sides of (5.3.1). This gives

$$h_\mu = \sum_{\lambda} \langle s_\lambda, h_\mu \rangle s_\lambda = \sum_{\lambda} \sum_{\rho} K_{\lambda\rho} \langle m_\rho, h_\mu \rangle s_\lambda = \sum_{\lambda \vdash n} K_{\lambda\mu} s_\lambda. \quad \square$$

Corollary 5.5.3.

$$e^{h_1} = \sum_{\lambda} \frac{f^\lambda}{|\lambda|!} s_\lambda.$$

Proof.

$$e^{h_1} = \sum_{n \geq 0} \frac{1}{n!} h_{1^n} = \sum_{n \geq 0} \frac{1}{n!} \sum_{\lambda \vdash n} K_{\lambda, 1^n} s_\lambda$$

But by definition, $K_{\lambda, 1^n} = \#\text{SYT}(\lambda) = f^\lambda$. □

Corollary 5.5.4 (Exponential specialization). *There is a ring homomorphism $\text{ex} : \Lambda \rightarrow \mathbb{Q}$ satisfying*

$$\text{ex}(s_\lambda) = \frac{f^\lambda}{|\lambda|!}.$$

Proof. Define $\text{ex} : \Lambda(x) \rightarrow \mathbb{Q}$ by $\text{ex}(h_n(x)) = \frac{1}{n!}$. Then

$$\sum_{\lambda} s_\lambda(x) s_\lambda(y) = \prod_{i,j \geq 1} (1 - x_i y_j)^{-1} = \prod_{j \geq 1} \sum_{n \geq 0} h_n(x) y_j^n$$

Applying ex to both sides gives

$$\begin{aligned} \sum_{\lambda} \text{ex}(s_\lambda(x)) s_\lambda(y) &= \prod_{j \geq 1} \sum_{n \geq 0} \frac{1}{n!} y_j^n \\ &= \prod_{j \geq 1} e^{y_j} = e^{h_1(y)} \\ &= \sum_{\lambda} \frac{f^\lambda}{|\lambda|!} s_\lambda(y). \end{aligned}$$

Comparing coefficients of $s_\lambda(y)$ we find $\text{ex}(s_\lambda(x)) = \frac{f^\lambda}{|\lambda|!}$. □

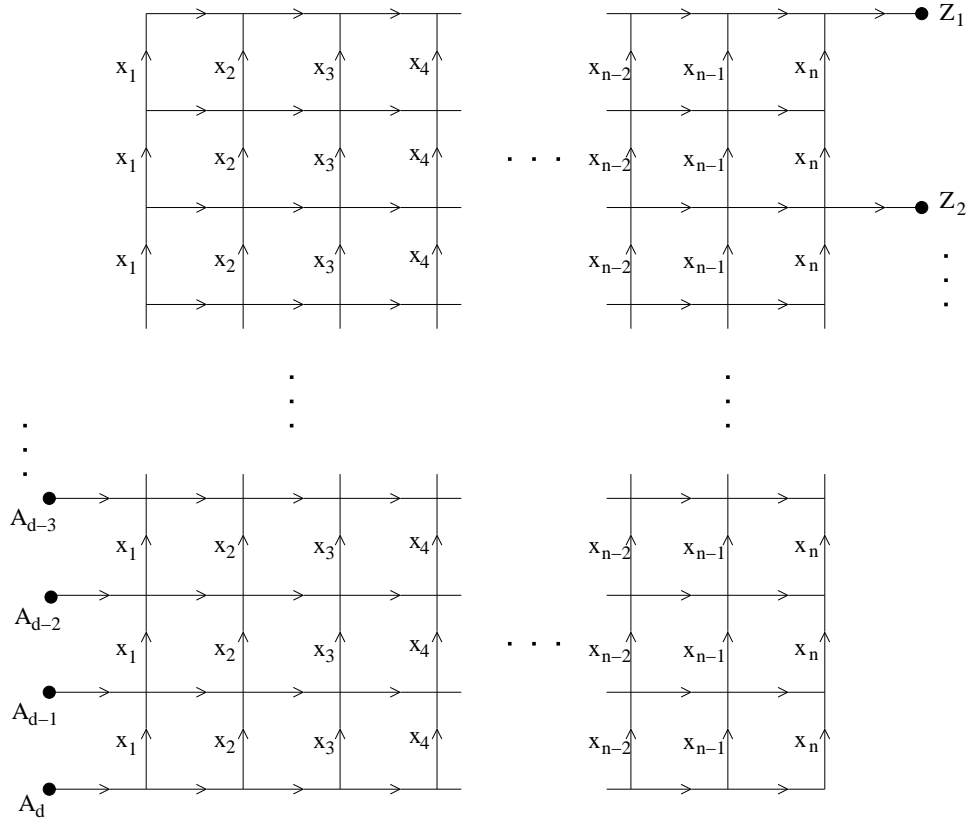


Figure 5.6.1: The grid used in the proof of the Jacobi-Trudi formula.

5.6 The Jacobi-Trudi formula

It is also possible to write the Schur functions explicitly in terms of the complete symmetric functions.

Theorem 5.6.1 (Jacobi-Trudi formula). *If $\lambda = (\lambda_1, \dots, \lambda_d)$, then*

$$s_\lambda = \det(h_{\lambda_i - i + j})_{i,j=1,\dots,d},$$

where by convention $h_0 = 1$ and $h_{-1} = h_{-2} = \dots = 0$.

Example 5.6.2. The Jacobi-Trudi formula states that

$$s_{21} = \det \begin{pmatrix} h_2 & h_3 \\ h_0 & h_1 \end{pmatrix} = h_{21} - h_3.$$

Which we can verify by working in $\Lambda^{(3)}$.

Note that the parts of λ appear on the diagonal of the matrix $(h_{\lambda_i - i + j})$.

Proof. We use the Gessel-Viennot Theorem on the grid shown in Figure 5.6.1, working in the ring $\Lambda^{(n)}$. The weight labels for the vertical edges are shown on the grid; the horizontal

edges have weight 1. The vertex Z_i is exactly λ_i rows above A_i , and A_i is one row above A_{i+1} .

First, observe that M_{ij} , the sum of all weights of paths from A_i to Z_j , is $h_{\lambda_j-j+i}(x_1, \dots, x_n)$. By Gessel-Viennot, the generating function for tuples of vertex-disjoint paths is therefore given by $\det(M_{ij}) = \det(h_{\lambda_j-j+i})$.

Now, given a collection of paths $P_i : A_i \rightarrow Z_i$, we produce a filling of the diagram of λ putting $a_1, a_2, \dots, a_{\lambda_i}$ in row i (in increasing order), if the vertical steps in P_i have weights $x_{a_1}, x_{a_2}, \dots, x_{a_{\lambda_i}}$. It is not hard to see that this filling is a SSYT (i.e. columns are strictly increasing) if and only if the paths P_i are vertex-disjoint paths, and this correspondence is bijective. Moreover, under this correspondence the weight of (P_1, \dots, P_n) is $x_1^{c_1} \dots x_n^{c_n}$ where (c_1, \dots, c_n) is the content of the corresponding tableau. It follows that the generating function for non-intersecting paths is $s_\lambda(x_1, \dots, x_n)$. \square

This proof has a number of great features. For example, it can be easily adapted to give a formula for s_λ in terms of the elementary symmetric functions. It can also be used to give a similar determinantal formula for *skew Schur functions*, $s_{\lambda/\mu}$, the generating function for SSYT(λ/μ). These can be found in Stanley's "Enumerative Combinatorics, Vol 2".

Corollary 5.6.3. *If $\lambda = (\lambda_1, \dots, \lambda_d)$,*

$$\frac{f^\lambda}{|\lambda|!} = \det \left(\frac{1}{(\lambda_i - i + j)!} \right)_{i,j=1,\dots,d},$$

where by convention $\frac{1}{k!} = 0$ for $k < 0$.

This is immediate from applying the exponential specialization ex to both sides of the Jacobi-Trudi formula. One can evaluate this determinant, using the Vandermonde identity: if $P_i(z)$ is a monic polynomial of degree $d - i$, then

$$\det (P_j(z_i))_{i,j=1,\dots,d} = \prod_{i < j} (z_i - z_j).$$

If we multiply each row of $\det \left(\frac{1}{(\lambda_i - i + j)!} \right)$ by $(\lambda_i - i + d)!$ we get $\det \left(\frac{(\lambda_i - i + d)!}{(\lambda_i - i + j)!} \right)$, which is of Vandermonde type, with $z_i = \lambda_i - i$. Thus we obtain

$$\frac{f^\lambda}{|\lambda|!} = \frac{\prod_{i < j} ((\lambda_i - i) - (\lambda_j - j))}{\prod_i (\lambda_i - i + d)!}. \quad (5.6.1)$$

This is called the *degree* formula for f^λ .

Another famous formula for f^λ is called the *hook* formula, which states

$$\frac{f^\lambda}{|\lambda|!} = \frac{1}{\prod_{\alpha \in \lambda} h(\alpha)}, \quad (5.6.2)$$

where $\alpha \in \lambda$ means that α is a cell in the Ferrers diagram of λ . The value of $h(\alpha)$, often called the "hook-length" for cell α , is the number of cells strictly to the right of and in the same row as α , plus the number of cells strictly below and in the same column as α , plus 1 (for the cell α itself). The term "hook-length" is used because the set of squares counted by $h(\alpha)$ form a "hook" shape. (See Figure:5.6.2.)




Figure 5.6.2: The hook length of α is 7.

Example 5.6.4. If $\lambda = (3, 2)$ the degree formula (5.6.1) gives

$$f^\lambda = \frac{5!((3-1) - (2-2))}{(3-1+2)!(2-2+2)!} = 5.$$

The hook formula (5.6.2) gives

$$f^\lambda = \frac{5!}{4 \cdot 3 \cdot 2 \cdot 1 \cdot 1} = 5,$$

in agreement with the degree formula.

Theorem 5.6.5 (Classical definition of Schur function). *If λ has at most d parts, then in $\Lambda^{(d)}$ the Schur function $s_\lambda(x_1, \dots, x_d)$ is given by*

$$s_\lambda = \frac{\det(x_j^{\lambda_i - i + d})_{i,j=1,\dots,d}}{\det(x_j^{-i+d})_{i,j=1,\dots,d}}.$$

Proof. (Omitted in class.) Working in $\Lambda^{(d)}$, consider

$$-\frac{1}{t}H(t)\frac{\partial}{\partial x_j}E(-t) = -\frac{1}{t}H(t)\frac{\partial}{\partial x_j}\prod_{k=0}^d(1 - x_k t) = -\frac{1}{t}H(t)E(-t)\frac{-t}{1 - x_j t} = -\frac{1}{1 - x_j t}.$$

Expanding both sides gives equating coefficients of $t^{\lambda_i - i + d}$ on both sides gives

$$\begin{aligned} [t^{\lambda_i - i + d}] \sum_{l \geq 0} h_l t^l \sum_{m=0}^{d-1} \frac{\partial}{\partial x_j} e_{m+1} (-1)^m t^m &= [t^{\lambda_i - i + d}] \frac{1}{1 - x_j t} \\ \sum_{k=1}^d (-1)^{d-k} \frac{\partial}{\partial x_j} e_{d-k+1} h_{\lambda_i - i + k} &= x_j^{\lambda_i - i + d} \end{aligned}$$

for all $i, j = 1, \dots, d$. These equations, together, can be written as the single matrix equation

$$(h_{\lambda_i - i + j})_{i,j=1,\dots,d} A = (x_j^{\lambda_i - i + d})_{i,j=1,\dots,d},$$

where

$$A = \left((-1)^{d-i} \frac{\partial}{\partial x_j} e_{d-i+1} \right)_{i,j=1,\dots,d}.$$

By Jacobi-Trudi, taking determinants gives us

$$s_\lambda \det A = \det (x_j^{\lambda_i - i + d})_{i,j=1,\dots,d}, \quad (5.6.3)$$

In the case where $\lambda = \epsilon$ (i.e., $\lambda_1 = \dots = \lambda_d = 0$), (5.6.3) yields

$$1 \det A = \det (x_j^{-i+d})_{i,j=1,\dots,d}. \quad (5.6.4)$$

Dividing (5.6.3) by (5.6.4) gives the result. \square

5.7 Skew Schur functions

We'll now consider skew Schur functions, which are the generating functions for SSYT of skew shape:

$$s_{\lambda/\mu} = \sum_{T \in \text{SSYT}(\lambda/\mu)} x_1^{c_1(T)} x_2^{c_2(T)} \dots$$

where $(c_1(T), c_2(T), \dots)$ is the content of T . The same proof that shows Schur functions are symmetric also works for skew Schur functions. It is convenient to adopt the convention that $s_{\lambda/\mu} = 0$ if $\mu \not\subseteq \lambda$.

Lemma 5.7.1. *For any partition μ , we have*

$$\sum_{\nu} s_\mu(x) s_\nu(x) s_\nu(y) = \sum_{\lambda} s_\lambda(x) s_{\lambda/\mu}(y). \quad (5.7.1)$$

Proof. Since

$$\sum_{\nu} s_\mu(x) s_\nu(x) s_\nu(y) = s_\mu(x) \prod_{i,j \geq 1} (1 - x_i y_j)^{-1},$$

the left hand side is the generating function for pairs $(T, \{(a_1, b_1), \dots, (a_k, b_k)\})$, where $T \in \text{SSYT}(\mu)$ and $\{(a_1, b_1), \dots, (a_k, b_k)\}$ is a multiset of pairs of positive integers. The right hand side is

$$\sum_{\lambda \supseteq \mu} s_\lambda(x) s_{\lambda/\mu}(y),$$

which is the generating function for pairs (P, Q) of semistandard Young tableaux where P has shape λ and Q has skew shape λ/μ for some $\lambda \supseteq \mu$.

We can prove the result by giving a content preserving bijection between these two sets. The bijection will be defined by a variation on the RSK correspondence, where instead of starting with $P_0 = Q_0 = \epsilon$, we begin with $P_0 = T$ and Q_0 empty. We define P_i and Q_i as before; if P_i has shape λ_i , then Q_i will have shape λ_i/μ . (See Figure 5.7.1 for an example.) This gives the desired bijection, by exactly the same argument as in the proof of Theorem 5.5.1. \square

Theorem 5.7.2. *For any $f \in \Lambda$, we have*

$$\langle s_\lambda, s_\mu f \rangle = \langle s_{\lambda/\mu}, f \rangle.$$




Figure 5.7.1: Example of the modified RSK correspondence, with $T = P_0$ as above, and multiset of pairs $\{(1, 6), (2, 2), (2, 6), (5, 1), (5, 4), (5, 4)\}$.

Proof. We will prove this for $f = s_\nu$, which is enough since the Schur functions form a basis for Λ . Taking the coefficient of $s_\nu(y)$ on both sides of (5.7.1) gives

$$s_\mu(x)s_\nu(x) = \sum_{\lambda \supseteq \mu} s_\lambda(x) \langle s_{\lambda/\mu}(y), s_\nu(y) \rangle_y,$$

and then taking the coefficient of $s_\lambda(x)$ gives

$$\langle s_\lambda(x), s_\mu(x)s_\nu(x) \rangle_x = \langle s_{\lambda/\mu}(y), s_\nu(y) \rangle_y,$$

as required □

We can write $s_{\lambda/\mu}$ in terms of the basis $\{m_\nu\}$

$$s_{\lambda/\mu} = \sum_{\nu} K_{\lambda/\mu, \nu} m_\nu. \tag{5.7.2}$$

From the definition of $s_{\lambda/\mu}$, $K_{\lambda/\mu;\nu}$ is the number of SSYT of shape λ/μ and content ν .

But now we can dualize this:

Corollary 5.7.3 (Iterated Pieri formula).

$$s_\mu h_\nu = \sum_\lambda K_{\lambda/\mu;\nu} s_\lambda$$

Proof. Apply $\sum_\lambda \langle \cdot, h_\nu \rangle s_\lambda$ to both sides of (5.7.2). This gives

$$\sum_\lambda \langle s_{\lambda/\mu}, h_\nu \rangle s_\lambda = \sum_\lambda \sum_\rho K_{\lambda/\mu;\rho} \langle m_\rho, h_\nu \rangle s_\lambda$$

Using Theorem 5.7.2, the left hand side simplifies to

$$\sum_\lambda \langle s_\lambda, s_\mu h_\nu \rangle s_\lambda = s_\mu h_\nu,$$

while the right simplifies to $\sum_\lambda K_{\lambda/\mu;\nu} s_\lambda$. □

Corollary 5.7.3 tells us how to multiply $s_\mu h_\nu$ and expand in the basis $\{s_\lambda\}$. The coefficients $K_{\lambda/\mu;\nu}$ that appear have a combinatorial meaning. Our next goal will be to give a combinatorial interpretation to the coefficients that appear in the product $s_\mu s_\nu$. This is where the crystal operators will come in to play.

A special case of this is when ν has only one part. In this case $h_\nu = h_k = s_k$. We say that a skew partition λ/μ is a horizontal strip, if it does not have two boxes in any column. We write $\mu \xrightarrow{k} \lambda$ to mean that λ/μ is a horizontal strip with k boxes.

Corollary 5.7.4 (Pieri formula).

$$s_\mu s_k = s_\mu h_k = \sum_{\mu \xrightarrow{k} \lambda} s_\lambda.$$

Proof. A SSYT of content $(k, 0, 0, \dots)$ is a tableau where all entries are 1. There is exactly one such SSYT of shape λ/μ , if λ/μ is a horizontal strip with k boxes, and since we cannot have two 1's in the same column of a SSYT, there are none if λ/μ is not a horizontal strip with k boxes. Thus $K_{\lambda/\mu;k} = 1$ if $\mu \xrightarrow{k} \lambda$ and $K_{\lambda/\mu;k} = 0$ otherwise. □

It is also possible to prove the Pieri formula directly, using the tableau product defined by row-insertion: one can show that the the tableau product $U * T$ defines a bijection $\text{SSYT}(\mu) \times \text{SSYT}(k) \rightarrow \coprod_{\mu \xrightarrow{k} \lambda} \text{SSYT}(\lambda)$. From the Pieri rule, the Iterated Pieri rule (Corollary 5.7.3) can be proved by induction; now the Proof of Corollary 5.7.3 that we gave can be run in reverse to give an alternate proof that $\{s_\lambda\}$ is an orthonormal basis for Λ . This alternate line of argument is quite similar to (and roughly the same amount of work as the approach) we took here.

Exercise 5.7.5. Work out the details for these alternate arguments.

5.8 The Littlewood-Richardson rule

Now define the set of *Littlewood-Richardson tableaux* (or LR tableaux) of shape λ/μ and content ν to be

$$\mathcal{LR}_{\lambda/\mu;\nu} = \{T \in \text{SSYT}(\lambda/\mu) \mid E_a(T) = \emptyset \text{ for all } a \geq 1\},$$

Here E_a is the crystal raising operator. The numbers $c_{\mu,\nu}^\lambda = \#\mathcal{LR}_{\lambda/\mu;\nu}$ are called *Littlewood-Richardson numbers*.

Theorem 5.8.1 (Littlewood-Richardson rule).

$$s_\mu s_\nu = \sum_{\lambda} c_{\mu,\nu}^\lambda s_\lambda$$

To prove this, we define a (labelled) directed graph $\Gamma(\lambda/\mu)$ with vertex set $\text{SSYT}(\lambda/\mu)$ by the following rule: if $U = E_a(T)$ for some $a \geq 1$, form a directed edge $T \rightarrow U$ (labelled by a). This graph is called the *crystal* on $\text{SSYT}(\lambda/\mu)$. The LR tableaux are precisely those vertices which have outdegree 0 in this graph.

In general crystals are not connected graphs; however, for straight shapes, they are. First note that $\Gamma(\nu)$ is an acyclic graph, since E_a causes one entry of the input tableau to decrease by 1 (while all other entries stay the same). Starting from any tableau in $\text{SSYT}(\nu)$, we can therefore follow a directed path in the graph until we arrive at a LR tableau. We show that the crystal $\Gamma(\nu)$ is connected by proving that this LR tableau is unique.

Lemma 5.8.2. *If ν is a partition, there is a unique LR tableau in $\text{SSYT}(\nu)$; it has the property that every entry in row i equals i . In particular, the crystal $\Gamma(\nu)$ is connected.*

Proof. Now suppose $T \in \text{SSYT}(\nu)$ is a LR tableau. We claim that the rightmost entry of the first row of T is 1. To see this, observe that this entry is the first element in the reading word; if this entry were $a > 1$, then this a survives the cancellation process (it has nothing before it to pair with), and so $E_{a-1}(T) \neq \emptyset$. It follows that every entry in the first row of T equals 1. By an inductive argument, we find that the every entry in row i equals i . \square

Lemma 5.8.3. *If C is any connected component of $\Gamma(\lambda/\mu)$, then C is isomorphic to $\Gamma(\nu_C)$ for some partition ν_C . Moreover, the isomorphism preserves content.*

Proof. Given any $T \in \text{SSYT}(\lambda/\mu)$ we can perform a sequence of slides to produce a straight shaped SSYT of some shape ν . Since sliding is reversible, and commutes with E_a for all a , this sequence of slides is an isomorphism from the component of T to $\Gamma(\nu)$. Since sliding does not change content, this isomorphism is content preserving. \square

Proof of the Littlewood-Richardson rule. The generating function for a component C of $\Gamma(\lambda/\mu)$ (where the weight is content) is s_{ν_C} ; hence, summing over all components we have

$$s_{\lambda/\mu} = \sum_C s_{\nu_C}.$$

But there is a unique LR tableau in C and its content is ν_C , the content of the unique LR tableau in $\Gamma(\nu_C)$. Hence we see that

$$\sum_C s_{\nu_C} = \sum_{\nu} c_{\mu,\nu}^{\lambda} s_{\nu}.$$

Thus we have $s_{\lambda/\mu} = \sum_{\nu} c_{\mu,\nu}^{\lambda} s_{\nu}$. and the result follows by Theorem 5.7.2 □

This argument shows ν_C is determined by C , from which we see that the isomorphism in Lemma 5.8.3 is actually unique. This observation allows us to prove another non-trivial theorem.

Theorem 5.8.4. *Let T be a SSYT of skew shape. The rectification of T , obtained by performing a sequence of slides until we have a SSYT of straight shape, is well defined (i.e. does not depend on the choices.)*

Proof. Any such sequence of slides defines the unique isomorphism from the crystal component C containing T to $\Gamma(\nu_C)$, and so is independent of the particular sequence. □

There are a number of alternate descriptions of LR tableaux that don't mention the crystal operators:

1. T is a LR tableau iff the word of T , $w(T) = w_1, \dots, w_n$, has the property that in any initial segment w_1, \dots, w_k , $1 \leq k \leq n$, the number of 1's \geq the number of 2's \geq the number of 3's $\geq \dots$
2. T is a LR tableau iff the tableau

$$\epsilon \leftarrow w_n \leftarrow w_{n-1} \leftarrow \dots \leftarrow w_1$$

has the property that every entry in row i equals i , for all $i \geq 1$.

3. T is a LR tableau with content ν iff there exists $U \in \text{SSYT}(\nu)$ with the property that the number of i 's in row j of T equals the number of j 's in row i of U , for all $i, j \geq 1$.

Exercise 5.8.5. Show that these really are alternate descriptions of LR tableaux.

From the Littlewood-Richardson rule, we can deduce several things. We call λ/μ a vertical strip if it does not have two boxes in any row. Equivalently, λ^t/μ^t is a horizontal strip.

Corollary 5.8.6 (Dual Pieri Formula).

$$s_{\mu} s_{1^k} = s_{\mu} e_k = \sum_{\mu^t \xrightarrow{k} \lambda^t} s_{\lambda}.$$

Proof. A Littlewood-Richardson tableau T of shape λ/μ and content 1^k , must have reading word $1, 2, 3, \dots, k$. Therefore, it is unique if it exists, and it cannot have the two entries in the same row (since they would be out of order). □

Corollary 5.8.7 (Iterated Dual Pieri Formula).

$$e_\mu = \sum_{\lambda} K_{\lambda\mu} s_{\lambda^t}$$

Proof. Apply the Dual Pieri formula inductively. □

Corollary 5.8.8. For every partition λ , $\omega(s_\lambda) = s_{\lambda^t}$.

Proof. We have

$$\sum_{\lambda} K_{\lambda,\mu} s_{\lambda^t} = e_\mu = \omega(h_\mu) = \sum_{\lambda} K_{\lambda,\mu} \omega(s_\lambda).$$

The result follows from the fact that $K_{\lambda,\mu}$ is an invertible matrix. □

Corollary 5.8.9 (Dual Jacobi-Trudi formula). If $\lambda = (\lambda_1, \dots, \lambda_d)$, then

$$s_{\lambda^t} = \det(e_{\lambda_i - i + j})_{i,j=1,\dots,d},$$

where by convention $e_0 = 1$ and $e_{-1} = h_{-2} = \dots = 0$.

Proof. Apply ω to both sides of the Jacobi-Trudi formula. □

5.9 Plane Partitions

One application of symmetric function theory is to enumerate *plane partitions*. These are two dimensional analogues of decreasing sequences of numbers. If λ is a partition, a plane partition with support λ is a filling Π of the boxes of λ such that the entries are weakly decreasing along rows and down columns. (Hence, if λ has only a single row, this is the same thing as a partition.) The size of Π is the sum of its entries, and is denoted $|\Pi|$.

We enumerate plane partitions, weighted by their sum, by producing a bijection with pairs of “reverse” semistandard Young tableaux. We’ll illustrate this by example. Suppose Π is the plane partition

5	5	5	4	3	1	1
4	4	3	1	1		
4	2	2	1			
2	2					
2	1					

Let λ be the entries on the main diagonal of Π ; in this case $\lambda = 542$. We will define two fillings P and Q of λ . Let α_i be the partition seen in Π , starting at the i^{th} diagonal entry, and reading to the right:

$$\alpha_1 = 5554311 \quad \alpha_2 = 4311 \quad \alpha_3 = 21.$$

Let β_i be the partition seen in Π , starting one box below the i^{th} diagonal entry, and reading downward (the last of these may be empty):

$$\beta_1 = 4422 \quad \beta_2 = 221 \quad \beta_3 = \epsilon.$$

We define P to be the filling of λ whose i^{th} row is $(\alpha_i)^t$. In this case,

$$P = \begin{array}{|c|c|c|c|c|} \hline 7 & 5 & 5 & 4 & 3 \\ \hline 4 & 2 & 2 & 1 & \\ \hline 2 & 1 & & & \\ \hline \end{array}$$

Similarly define Q to be the filling of λ whose i^{th} row is $(\beta_i)^t$. This produce an object with shape $\nu \subseteq \mu$; we put zeros in the empty boxes:

$$Q = \begin{array}{|c|c|c|c|c|} \hline 4 & 4 & 2 & 2 & 0 \\ \hline 3 & 2 & 0 & 0 & \\ \hline 0 & 0 & & & \\ \hline \end{array}$$

One should check the following facts:

1. The sum of the entries in P and Q together is $|\Pi|$.
2. The entries of P and Q will be weakly decreasing along rows, and weakly increasing down columns.
3. This map is a *bijection* between plane partitions and pairs (P, Q) of “reverse” SSYT of the same shape, where P has entries in $\mathbb{Z}_{>0}$ and Q has entries in $\mathbb{Z}_{\geq 0}$,

The generating function for the set of all such P of shape λ , weighted by sum of the entries, is $s_\lambda(x, x^2, x^3, \dots)$, while the corresponding generating function for the set of all Q is $s_\lambda(1, x, x^2, \dots)$. It follows that the generating function for plane partitions is

$$\begin{aligned} \sum_{\lambda} s_\lambda(x, x^2, x^3, \dots) s_\lambda(1, x, x^2, \dots) &= \prod_{i, j \geq 1} (1 - (x^i)(x^{j-1}))^{-1} \\ &= \prod_{k \geq 1} (1 - x^k)^{-k}, \end{aligned}$$

which is a classical result of MacMahon. (For the last equality, we use the fact that there are k pairs of $i, j \geq 1$ such that $i + j - 1 = k$.)

Chapter 6

Representations of the symmetric group

One important application of symmetric function theory is to the representation theory of the symmetric group S_n . This in turn has a number of combinatorial applications. In this chapter, we'll give an overview of the basics of representation theory, construct the representations of S_n , and compute their characters. We'll then apply this theory to the combinatorial problem of enumerating maps on a surface.

6.1 Definitions and examples

We begin by recalling some basic facts and notation about S_n . A permutation $\pi \in S_n$ is most commonly specified as the ordered list

$$\pi(1), \pi(2), \dots, \pi(n).$$

Permutations define a group, where multiplication is given by composition of functions. The identity element is the identity function

$$\text{id} = 1, 2, 3, \dots, n.$$

Another way to specify a permutation is in cycle notation. We write $(i_1 \ i_2 \ \dots \ i_k)$ to mean the permutation π that has $\pi(i_1) = i_2, \dots, \pi(i_k) = i_1$, and $\pi(j) = j$ for all $j \neq i_1, i_2, \dots, i_k$. Any permutation can be uniquely written as a product of disjoint cycles; the lengths of those cycles determine the conjugacy class of the permutation. A cycle with two element $(a \ b)$ is called a transposition, and will also be denoted r_{ab} .

In order to talk about representation theory of S_n , we need some general theory which will apply to any finite group G . A (finite dimensional, complex) *representation* of G consists of a finite dimensional complex vector space V and a group homomorphism $\rho_V : G \rightarrow \text{GL}(V)$. (A homomorphism is a map satisfying $\rho(1) = I_V$, and $\rho_V(gh) = \rho_V(g)\rho_V(h)$ for all $g, h \in G$; $\text{GL}(V)$ is the set of invertible linear operators on V .) It is common practice to abuse terminology and call V alone (without reference to ρ_V) the representation. It is also common to write gx to mean $\rho_V(g)x$, where $x \in V, g \in G$.

If V is a representation of G . A subspace $W \subset V$ is called a subrepresentation if W is $\rho_V(g)$ -invariant (i.e. $\rho_V(g)W \subset W$) for all $g \in G$. A subrepresentation $W \subset V$ is a representation of G (where $\rho_W(g) = \rho_V(g)|_W$ for all $g \in G$).

Before we get any deeper into the theory, it is worth looking at some examples of representations and subrepresentations for the group S_n .

Example 6.1.1. Let $V = \mathbb{C}^n$ with standard basis $\{e_1, \dots, e_n\}$, and for $\pi \in S_n$, let $\rho(\pi)$ be the permutation matrix defined by $\rho(\pi)e_i = e_{\pi(i)}$. Then V is a representation of S_n .

The subspace T spanned by $\{(1, \dots, 1)\}$ is a subrepresentation. This is a “trivial” representation, in that $\pi x = x$ for all $x \in T$, $\pi \in S_n$.

The subspace $\tilde{V} = \{(x_1, \dots, x_n) \mid \sum x_i = 0\}$ is also a subrepresentation. Notice that T and \tilde{V} are complementary to each other, in the sense that $V = T \oplus \tilde{V}$.

Example 6.1.2. Let $A = \mathbb{C}^1$, and for $\pi \in S_n$ define $\rho_A(\pi) = \text{sgn}(\pi)$, thought of as a 1×1 matrix. Since $\text{sgn}(\pi\sigma) = \text{sgn}(\pi)\text{sgn}(\sigma)$ this is a representation. A is called the alternating representation of S_n and is different from the one dimensional representation T .

Example 6.1.3. Let $G = S_n$ and let $V = \mathbb{C}[x_1, \dots, x_n]_k$ be the vector space of homogeneous polynomials of degree k . Let $\rho_V(\pi)f(x_1, \dots, x_n) = f(x_{\pi(1)}, \dots, x_{\pi(n)})$. Then V is a representation of S_n .

The subspace $\Lambda_k^{(n)}$ of symmetric polynomials of degree k in n variables, is a subrepresentation of V . Another example of a subrepresentation is the span of the monomials x_1^k, \dots, x_n^k .

Example 6.1.4. Let \mathcal{A} be a species. Define $\mathbb{C}\mathcal{A}_n$ to be a vector space whose basis is $\mathcal{A}_{[n]}$. Thus the elements of $\mathbb{C}\mathcal{A}_n$ are formal linear combinations of elements of \mathcal{A} -structures on $\{1, \dots, n\}$. If $\pi \in S_n$ we define $\rho_{\mathbb{C}\mathcal{A}_n}$ by setting $\rho_{\mathbb{C}\mathcal{A}_n}(\pi)(A) = \pi_*(A)$ for $A \in \mathcal{A}_{[n]}$, and extending this to a linear transformation. This makes $\mathbb{C}\mathcal{A}_n$ a representation of S_n .

For example, for the species \mathcal{E}^\bullet , an \mathcal{E}^\bullet -structure on $[n]$ is a pair $e_i = ([n], i)$ where $i \in [n]$. Thus $\mathbb{C}\mathcal{E}^\bullet$ is the n -dimensional vector space spanned by $\{e_1, \dots, e_n\}$. If $\pi \in S_n$ then $\pi e_i = \pi_*(e_i) = e_{\pi(i)}$. So $\mathbb{C}\mathcal{E}^\bullet$ is the same as the representation V in Example 6.1.1. The representation \tilde{V} in Example 6.1.1, however is not of this form; neither is A in Example 6.1.2

Example 6.1.5. We’ll be most interested in a special case of this last example, where the species is

$$\mathcal{M}^\lambda = \mathcal{E}_{\lambda_1} * \mathcal{E}_{\lambda_2} * \dots * \mathcal{E}_{\lambda_d}.$$

Here, $\lambda = (\lambda_1, \dots, \lambda_d)$ is a partition of n . In this case the representation $\mathbb{C}\mathcal{M}_{[n]}^\lambda$ is given the name M^λ . An element of $\mathcal{M}_{[n]}^\lambda$ is a set composition (S_1, \dots, S_d) of $[n]$ where $\#S_i = \lambda_i$. We will represent such an element by a tableau-like diagram called a *tabloid*, wherein the elements of S_i are listed in i^{th} row. Often they are written in increasing numerical order, but that is simply convention; to indicate that the order of the elements within each row doesn’t matter we leave out the vertical line separating the boxes. For example,

$$\mathfrak{u} = \begin{array}{|c|c|c|c|} \hline 3 & 5 & 6 & 9 \\ \hline 1 & 4 & 8 & \\ \hline 2 & & & \\ \hline 7 & & & \\ \hline \end{array}$$

is a tabloid of shape 4311; if $\pi = 283941756$, then

$$\pi\mathfrak{U} = \begin{array}{|c|c|c|c|} \hline 3 & 4 & 1 & 6 \\ \hline 2 & 9 & 5 & \\ \hline 8 & & & \\ \hline 7 & & & \\ \hline \end{array} = \begin{array}{|c|c|c|c|} \hline 1 & 3 & 4 & 6 \\ \hline 2 & 5 & 9 & \\ \hline 8 & & & \\ \hline 7 & & & \\ \hline \end{array}$$

6.2 A crash course in representation theory

Proposition 6.2.1. *Let V be a representation of G . There is a Hermitian positive definite inner product $H_V(\cdot, \cdot)$ on V such that for all $x, y \in V, g \in G$,*

$$H_V(gx, gy) = H_V(x, y).$$

Such an inner product is called a G -invariant inner product.

Proof. Let $\tilde{H}(x, y)$ be any positive definite inner product on V . It is easy to check that $H_V(x, y) = \sum_{g \in G} \tilde{H}(gx, gy)$ has the desired properties. \square

Proposition 6.2.2. *Let V be a representation of G . If $W \subset V$ is a subrepresentation then there is another subrepresentation $W' \subset V$ such that $V = W' \oplus W$.*

Proof. Define $W' = W^\perp$, the orthogonal complement to W under a G -invariant inner product. Then $V = W \oplus W'$. Then if $y \in W'$ then $H_V(y, g^{-1}w) = H_V(gy, w)$ for all $w \in W$. so $gy \in W'$; hence W' is a subrepresentation. \square

We call a representation V reducible if it has a proper non-zero subrepresentation W , $\{0\} \neq W \subsetneq V$. We call V *irreducible* if it is not reducible. The representations T and \tilde{V} from Example 6.1.1, and A from Example 6.1.2 are examples of irreducible representations. The phrase “irreducible representation” is often abbreviated as “irrep”.

Corollary 6.2.3. *Every representation of V can be decomposed as $V = V_1 \oplus \cdots \oplus V_k$, where V_i is an irreducible representation.*

If V is a representation of G , then $V^G = \{x \in V \mid gx = x\}$ is a subrepresentation, called the space of invariant vectors. Note that $\rho_V g$ maps every group element to the identity. We have already seen some examples. In Example 6.1.1 the subrepresentation $T \subset \mathbb{C}^n$ is the space of invariant vectors; in Example 6.1.3, $\Lambda_k^{(n)} = (\mathbb{C}[x_1, \dots, x_n]_k)^{S_n}$.

Proposition 6.2.4. *The operator*

$$P_V = \frac{1}{|G|} \sum_{g \in G} \rho_V(g)$$

is a projection from V onto V^G .

Proof. One needs to check that if $x \in V$ then $P_V x \in V^G$, and if $x \in V^G$ then $P_V x = x$; both are easy. \square

If V and W are representations of G , we define $\text{Hom}(V, W)$ to be the space of linear transformations from V to W . Then $\text{Hom}(V, W)$ is a representation of G as follows: for $\phi \in \text{Hom}(V, W)$ define $g\phi : V \rightarrow W$ by

$$(g\phi)(x) = g\phi(g^{-1}x) \quad x \in V.$$

We say that a linear map $\phi \in \text{Hom}(V, W)$ is a G -homomorphism if $\phi \in \text{Hom}(V, W)^G$. Equivalently, $\phi(gx) = g(\phi(x))$ for all $g \in G, x \in V$. If ϕ is a G -isomorphism, we say that V and W are isomorphic representations.

Lemma 6.2.5 (Schur's lemma). *Let V and W be irreducible representations. If V and W are isomorphic G -representations, then $\dim \text{Hom}(V, W)^G = 1$. Otherwise $\dim \text{Hom}(V, W)^G = 0$.*

Proof. First we show that any $\phi \in \text{Hom}(V, W)^G$ is either the zero map or an isomorphism. Note that the image of ϕ and the kernel of ϕ are subrepresentations of V and W respectively. Since V is irreducible, $\ker \phi = \{0\}$ or $\ker \phi = V$. Similarly, $\text{im } \phi = 0$ or $\text{im } \phi = W$. Since $\dim(\ker \phi) + \dim(\text{im } \phi) = \dim V$, the only possibilities are $\ker \phi = \{0\}$, $\text{im } \phi = W$, in which case ϕ is an isomorphism, or $\ker \phi = V$, in which case ϕ is the zero map.

If V and W are not isomorphic, this shows that the only map in $\text{Hom}(V, W)^G$ is the zero map.

On the other hand, if V and W are isomorphic, there is a G -isomorphism $\phi \in \text{Hom}(V, W)^G$, so $\dim \text{Hom}(V, W)^G \geq 1$. To show the dimension is exactly one, suppose ϕ, ψ are both G -isomorphisms. Let x be any eigenvector of $\psi^{-1}\phi$, with eigenvalue λ . Then $\phi - \lambda\psi \in \text{Hom}(V, W)^G$, and $(\phi - \lambda\psi)x = 0$. As $\phi - \lambda\psi$ is not an isomorphism, it must be the zero map; hence $\phi = \lambda\psi$. \square

The character of a representation V is the function $\chi_V : G \rightarrow \mathbb{C}$ defined by

$$\chi_V(g) = \text{tr } \rho_V(g),$$

the trace of $\rho_V(g)$. Character theory is the main tool used to study representations (similar to the way the characteristic polynomial is used to study linear operators). We begin with three easy, but important observations:

1. Isomorphic representations have the same character. To see this, note that if $\phi : V \rightarrow W$ is a G -isomorphism, then $\rho_W(g) = \phi\rho_V(g)\phi^{-1}$, which have the same trace. (Eventually we'll show that the converse is also true: if two representations have the same character, then they are isomorphic.)
2. If $V = W \oplus W'$, where W and W' are subrepresentations, then $\chi_V = \chi_W + \chi_{W'}$. To see this, simply note that $\rho_V(g) = \rho_W(g) \oplus \rho_{W'}(g)$, is block diagonal for all $g \in G$. The trace of a block diagonal operator is the sum of traces of the blocks. For this reason, our main goal will be to understand the characters of irreducible representations—every other character is a sum of these.
3. The character takes the same value on conjugate elements, i.e. $\chi_V(g) = \chi_V(hgh^{-1})$. To see this, note that $\rho_V(g)$ and $\rho_V(hgh^{-1}) = \rho_V(h)\rho_V(g)\rho_V(h)^{-1}$ have the same trace. Functions which this property are called class functions. Thus the dimension of the subspace of functions from $G \rightarrow \mathbb{C}$ spanned by characters is at most the number of conjugacy classes of G . (In fact it is always equal, but we won't prove this in general.)

6.3 Character theory for S_n

For $G = S_n$ the conjugacy classes are the elements with the same cycle type. That is, π and σ are conjugate iff π and σ have the same number of cycles of length k for all k . We can represent the cycle type of π by a partition $\mu \vdash n$, where the number of parts of size k in μ is the number of cycles of length k in π . We write \mathcal{C}_μ for the conjugacy class of all permutations in S_n with cycle type μ . If $\mu \vdash n$, we'll write $\chi_V(\mu)$ to mean $\chi_V(\pi)$ for any $\pi \in \mathcal{C}_\mu$.

Example 6.3.1. Let $V = \mathbb{C}^n$ be the representation of S_n defined in Example 6.1.1. Then $\chi_V(\pi)$ is the number of 1's on the diagonal of the permutation matrix $\rho_V(\pi)$, which is the number of fixed points of π . Note that permutations with the same cycle type have the same number of fixed points. If $\mu \vdash n$, then $\chi_V(\mu)$ is the number of parts of size 1 in μ .

If T is the subrepresentation spanned by $\{1, \dots, 1\}$, then $\chi_T(\pi) = 1$ for all $\pi \in S_n$. It follows that if \tilde{V} is the complement to T , then $\chi_{\tilde{V}}(\pi) = \#\text{fixed points of } \pi - 1$.

More generally, for any species \mathcal{A} , the character $\chi_{\mathcal{C}\mathcal{A}_n}(\pi)$ is the number of \mathcal{A} -structures on $[n]$ that are fixed by π_* . The reasoning is the same: $\rho_{\mathcal{C}\mathcal{A}_n}(\pi)$ is represented by a permutation matrix and each fixed \mathcal{A} -structure gives a 1 on the diagonal. In particular, we can calculate the character of the representation M^λ . We denote this character by ξ^λ .

Proposition 6.3.2. *The character ξ^λ of the representation M^λ is given by*

$$\xi^\lambda(\mu) = \langle h_\lambda, p_\mu \rangle.$$

Proof. Let π be a permutation of cycle type $\mu = (\mu_1, \dots, \mu_k)$. Then a tabloid \mathfrak{U} of shape λ is fixed by π_* if and only if π_* preserves each row of \mathfrak{U} . This means that each cycle of π must be contained within some row of \mathfrak{U} , or equivalently that each row of \mathfrak{U} must be a union of cycles of π . Thus each such \mathfrak{U} gives a function $f : [n] \rightarrow \mathbb{Z}_{>0}$, defined by $f(j) = i$ if the j lies in the i^{th} row of \mathfrak{U} , satisfying $f \circ \pi = f$ and $\lambda_i = \#\{j \mid f(j) = i\}$. Conversely each such function comes from a tabloid of shape λ . The number of these functions is given by $[x_1^{\lambda_1} \cdots x_d^{\lambda_d}] p_\mu$, and therefore this is also the number of π_* -fixed tabloids of shape λ . Finally, using the fact that h_λ and m_λ dual bases under the Hall inner product on symmetric functions, this can be rewritten as $\langle h_\lambda, p_\mu \rangle$. \square

From these examples, one might get the impression that every character of G is an integer valued function. As we'll see, this is actually the case if $G = S_n$, but it is certainly not true for all groups. In general characters are not even real valued functions (try, for example, computing the characters for the three element cyclic group). The fact that the characters of S_n are real valued functions can be seen from the following proposition and the fact that π and π^{-1} have the same cycle type.

Proposition 6.3.3. *If V is any representation of G , then $\overline{\chi_V(g)} = \chi_V(g^{-1})$.*

The bar here denotes complex conjugation.

Proof. Since $H_V(g^{-1}x, y) = H_V(x, gy)$ for all $x, y \in V$, $\rho_V(g^{-1})$ is the Hermitian adjoint of $\rho_V(g)$. This implies that their traces are complex conjugates. \square

If V, W are representations, we define the inner product of their characters to be

$$\langle \chi_W, \chi_V \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_W(g) \overline{\chi_V(g)}.$$

(Note: this is a completely different inner product on a completely different vector space from the $H_V(\cdot, \cdot)$ defined earlier.)

For S_n the number of permutations with cycle type μ is $|\mathcal{C}_\mu| = \frac{n!}{z(\mu)}$. Using this, and the fact that the characters of S_n are real, this formula simplifies to

$$\langle \chi_W, \chi_V \rangle = \sum_{\mu \vdash n} \frac{1}{z(\mu)} \chi_W(\mu) \chi_V(\mu),$$

for representations V and W of S_n .

Lemma 6.3.4. *If V and W are representations of G , then the character of $\text{Hom}(V, W)$ is $\chi_{\text{Hom}(V, W)}(g) = \chi_W(g) \overline{\chi_V(g)}$.*

Proof. Let $H_V(\cdot, \cdot)$ and $H_W(\cdot, \cdot)$ be G -invariant inner products on V and W respectively, and let $\{e_i\}, \{f_j\}$ be orthonormal bases for V and W respectively.

Then the linear maps $\{\phi_{ij}\}$, defined by $\phi_{ij}(v) = H_V(v, e_i) f_j$, form a basis for $\text{Hom}(V, W)$; in particular, for any $\psi \in \text{Hom}(V, W)$,

$$\psi = \sum_{i,j} H_W(\psi(e_i), f_j) \phi_{ij}.$$

Thus when we expand $g\phi_{ij}$ in terms of this basis the coefficient of ϕ_{ij} is given by

$$H_W((g\phi_{ij})(e_i), f_j);$$

these numbers are the diagonal elements of the matrix representation of $g\phi_{ij}$. The character is therefore,

$$\begin{aligned} \chi_{\text{Hom}(V, W)}(g) &= \sum_{i,j} H_W((g\phi_{ij})(e_i), f_j) \\ &= \sum_{i,j} H_W(gH_V(g^{-1}e_i, e_i) f_j, f_j) \\ &= \sum_{i,j} H_W(gf_j, f_j) H_V(g^{-1}e_i, e_i) \\ &= \chi_W(g) \chi_V(g^{-1}) = \chi_W(g) \overline{\chi_V(g)}. \end{aligned} \quad \square$$

Theorem 6.3.5. *If V and W are representations of G , then*

$$\langle \chi_W, \chi_V \rangle = \dim \text{Hom}(V, W)^G.$$

In particular the characters for irreducible representations are orthonormal.

Proof. We have

$$\begin{aligned}
\langle \chi_W, \chi_V \rangle &= \frac{1}{|G|} \sum_{g \in G} \chi_W(g) \overline{\chi_V(g)} \\
&= \frac{1}{|G|} \sum_{g \in G} \chi_{\text{Hom}(V, W)}(g) \\
&= \text{tr} \left(\frac{1}{|G|} \sum_{g \in G} \rho_{\text{Hom}(V, W)}(g) \right) \\
&= \text{tr}(P_{\text{Hom}(V, W)}) \\
&= \dim \text{Hom}(V, W)^G
\end{aligned}$$

where the last equality is because the trace of a projection is the dimension of its image.

In particular, if V and W are irreducible, $\langle \chi_W, \chi_V \rangle = 1$ if V and W are isomorphic, and 0 otherwise. \square

Corollary 6.3.6. *Two representations of G are isomorphic if and only if they have the same character.*

Proof. Let V and W two representations of G . We already saw that if V is isomorphic to W then $\chi_V = \chi_W$. Conversely if $\chi_V = \chi_W$, consider the decompositions $V = V_1 \oplus \cdots \oplus V_k$ and $W = W_1 \oplus \cdots \oplus W_\ell$ into irreducible subrepresentations. We have $\sum_{i=1}^k \chi_{V_i} = \sum_{j=1}^{\ell} \chi_{W_j}$. Since the irreducible characters are linearly independent, this tells us W_1, \dots, W_ℓ is isomorphic to a permutation of V_1, \dots, V_k ; and so V is isomorphic to W . \square

Corollary 6.3.7. *The number of non-isomorphic irreducible representations of G is at most the number of conjugacy classes of G .*

Proof. The irreducible characters are linearly independent and belong to the vector space of class functions, whose dimension is the number of conjugacy classes of G . \square

The characters ξ^λ , $\lambda \vdash n$ are not orthonormal, which indicates that the representations M^λ are not irreducible. From Proposition 6.3.2 we have $\xi^\lambda(\mu) = 0$ if μ is lexicographically greater than λ , and $\xi^\lambda(\lambda) = \langle h_\lambda, p_\lambda \rangle \neq 0$. This implies that the characters ξ^λ , $\lambda \vdash n$ are linearly independent. This proves a number of things:

1. Every irreducible representation must be contained in some M^λ (up to isomorphism).
2. The number of non-isomorphic irreducible representations is the number of partitions of n .
3. Every character of S_n is a rational valued function.

6.4 Specht modules

Our next goal will construct an irreducible representation $S^\lambda \subset M^\lambda$ for each $\lambda \vdash n$. These will be all the irreducible representations of S_n . To do this, we will need to look at certain linear combinations of operators $\rho_{M^\lambda}(\pi)$, $\pi \in S_n$. It will be convenient to view this as a calculation in the *group algebra* of the symmetric group, which we now define.

Let $\mathbb{C}[S_n]$ to be the complex vector space with basis S_n ; thus the elements of $\mathbb{C}[S_n]$ are formal linear combinations $\sum_{\pi \in S_n} a_\pi \pi$ of elements of S_n . We define a product $\mathbb{C}[S_n] \times \mathbb{C}[S_n] \rightarrow \mathbb{C}[S_n]$ by linearly extending the product of S_n :

$$\left(\sum_{\pi \in S_n} a_\pi \pi \right) \left(\sum_{\sigma \in S_n} b_\sigma \sigma \right) = \sum_{\pi, \sigma} a_\pi b_\sigma (\pi \sigma).$$

This product makes $\mathbb{C}[S_n]$ a (non-commutative) ring.

Example 6.4.1. $\mathbb{C}[S_3]$ has basis $[123], [213], [132], [231], [312], [321]$. An example of a product in this ring is:

$$\begin{aligned} \left([213] + \frac{1}{3}[321] \right) \left(-[231] + 5[123] \right) &= -[213][231] + 5[213][123] - \frac{1}{3}[321][231] + \frac{5}{3}[321][123] \\ &= -[132] + 5[213] - \frac{1}{3}[213] + \frac{5}{3}[321] \\ &= -[132] + \frac{14}{3}[213] + \frac{5}{3}[321] \end{aligned}$$

Given a representation V of S_n , we can extend ρ_V to a linear map $\rho_V : \mathbb{C}[S_n] \rightarrow \text{End}(V)$, by putting

$$\rho_V \left(\sum_{\pi \in S_n} a_\pi \pi \right) = \sum_{\pi \in S_n} a_\pi \rho_V(\pi).$$

Note that the codomain of this map is now $\text{End}(V)$, the set of all linear operators on V , not just the invertible ones. This construction makes ρ_V a ring homomorphism. We'll continue to write ax for $\rho_V(a)(x)$, now for $a \in \mathbb{C}[S_n]$, $x \in V$.

An $[n]$ -filling of $\lambda \vdash n$ is a filling of the boxes in the Ferrers diagram of λ with entries $1, \dots, n$, each used once. Given an $[n]$ -filling T of $\lambda \vdash n$, we'll associate four objects.

1. A tabloid, denoted $\{T\} \in M^\lambda$, obtained by forgetting the order of the elements in each row of T .
2. A subgroup $C(T) \subset S_n$, which is the set of permutations π such that $\pi(i)$ is in the same column as i in T .
3. An element of $b_T \in \mathbb{C}[S_n]$, which is defined to be

$$b_T = \sum_{\pi \in C(T)} \text{sgn}(\pi) \pi.$$

These elements (and some other similar objects that we won't discuss) are known as *Young symmetrizers*.

4. A vector

$$v_T = b_T\{T\} \in M^\lambda.$$

Observe that applying $\pi \in S_n$ to each of the entries of T just changes the labels in all of these constructions, so $\pi v_T = v_{\pi T}$ for all $\pi \in S_n$.

Example 6.4.2. For the filling

$$T = \begin{array}{|c|c|c|} \hline 3 & 5 & 2 \\ \hline 1 & & \\ \hline 4 & & \\ \hline \end{array}$$

we have

$$\{T\} = \begin{array}{|c|c|c|} \hline 2 & 3 & 5 \\ \hline 1 & & \\ \hline 4 & & \\ \hline \end{array}$$

and $C(T)$ is the set of permutations of $\{1, 2, 3, 4, 5\}$ such that 2 and 5 are fixed points. Hence

$$b_T = [12345] - [32145] - [12435] + [32415] + [42135] - [42315],$$

and

$$v_T = \begin{array}{|c|c|c|} \hline 2 & 3 & 5 \\ \hline 1 & & \\ \hline 4 & & \\ \hline \end{array} - \begin{array}{|c|c|c|} \hline 1 & 2 & 5 \\ \hline 3 & & \\ \hline 4 & & \\ \hline \end{array} - \begin{array}{|c|c|c|} \hline 2 & 4 & 5 \\ \hline 1 & & \\ \hline 3 & & \\ \hline \end{array} \\ + \begin{array}{|c|c|c|} \hline 2 & 4 & 5 \\ \hline 3 & & \\ \hline 1 & & \\ \hline \end{array} + \begin{array}{|c|c|c|} \hline 1 & 2 & 5 \\ \hline 4 & & \\ \hline 3 & & \\ \hline \end{array} - \begin{array}{|c|c|c|} \hline 2 & 3 & 5 \\ \hline 4 & & \\ \hline 1 & & \\ \hline \end{array}.$$

Define the *Specht module* $S^\lambda \subset M^\lambda$ to be the vector space spanned by the vectors v_T where T is a $[n]$ -filling of λ . Since $\pi v_T = v_{\pi T} \in S^\lambda$ for all $\pi \in S_n$, S^λ is a subrepresentation.

Our next goals are to prove that S^λ is irreducible, and that S^λ and S^μ are non-isomorphic for $\lambda \neq \mu$. For this we need two lemmas, which study the action of the Young symmetrizers $b_T \in \mathbb{C}[S_n]$.

Lemma 6.4.3. *Let T be a $[n]$ -filling of $\lambda \vdash n$, let \mathfrak{U} be a tabloid of shape $\mu \vdash n$. If $a \neq b$ are in the same row of \mathfrak{U} and the same column of T then $b_T \mathfrak{U} = 0 \in M^\mu$.*

Proof. We define a sign reversing involution on the set $C(T)$, where the weight of a permutation $\pi \in C(T)$ is $\pi \mathfrak{U}$, and the index is $\ell(\pi)$. Since a and b are in the same column of T the transposition $r_{ab} = (a, b)$ is in $C(T)$. We define $\alpha(\pi) = \pi r_{ab}$, which is an involution. Since a and b are in the same row of \mathfrak{U} , we have $r_{ab} \mathfrak{U} = \mathfrak{U}$. We have $\text{ind}(\alpha(\pi))$ is odd iff $\text{ind}(\pi)$ is even, and

$$\text{WT}(\pi) = \pi \mathfrak{U} = \pi r_{ab} \mathfrak{U} = \text{WT}(\alpha(\pi)).$$

Thus α is a sign reversing involution and so $b_T \mathfrak{U} = \sum_{\pi \in C(T)} (-1)^{\text{ind}(T)} \text{WT}(T) = 0$. \square

Lemma 6.4.4. *If T is $[n]$ -filling of λ , then for any $x \in M^\lambda$, $b_T x$ is a multiple of v_T . Moreover $b_T v_T = |C(T)|v_T \neq 0$.*

Proof. To prove this, we show that $b_T \mathfrak{U}$ is a multiple of v_T for any tabloid \mathfrak{U} of shape λ . If there exist $a \neq b$ in the same row of \mathfrak{U} and the same column of T then $b_T \mathfrak{U} = 0$. Otherwise, we can find permutation $\sigma \in C(T)$ such that $\mathfrak{U} = \{\sigma(T)\}$. But then

$$\begin{aligned} b_T \mathfrak{U} &= \sum_{\pi \in C(T)} \operatorname{sgn}(\pi) \pi \mathfrak{U} \\ &= \sum_{\pi \in C(T)} \operatorname{sgn}(\pi) \pi \sigma \{T\} \\ &= \operatorname{sgn}(\sigma) \sum_{\pi \in C(T)} \operatorname{sgn}(\pi \sigma) \pi \sigma \{T\} \\ &= \operatorname{sgn}(\sigma) v_T. \end{aligned}$$

For the final statement, we have

$$\begin{aligned} b_T v_T &= \sum_{\pi' \in C(T)} \operatorname{sgn}(\pi') \pi' \sum_{\pi \in C(T)} \operatorname{sgn}(\pi) \pi \{T\} \\ &= \sum_{\pi', \pi \in C(T)} \operatorname{sgn}(\pi' \pi) \pi' \pi \{T\} \\ &= |C(T)| \sum_{\pi \in C(T)} \operatorname{sgn}(\pi) \pi \{T\} \\ &= |C(T)| v_T. \end{aligned} \quad \square$$

Theorem 6.4.5. S^λ is irreducible.

Proof. Suppose to the contrary that $S^\lambda = W \oplus W'$, where W and W' are proper subrepresentations. We can write v_T as $w + w'$ where $w \in W$ and W' in W' . Then $b_T w \in W$ and $b_T w' \in W'$. Since $0 \neq b_T v_T = b_T w + b_T w'$, we can assume without loss of generality that $b_T w \neq 0$, and hence is a non-zero multiple of v_T . This tells us $v_T \in W$. But then we have $\pi v_T = v_{\pi T} \in W$ for all $\pi \in S_n$, and these vectors span S^λ ; hence $W = S^\lambda$ which is a contradiction. \square

Theorem 6.4.6. *If μ is lexicographically greater than λ then $\dim \operatorname{Hom}(S^\lambda, M^\mu)^{S_n} = 0$.*

In particular there are no non-zero S_n -homomorphisms between S^λ and S^μ if μ is lexicographically greater than λ , which tells us that S^λ and S^μ are non-isomorphic. Thus we have found all the irreducible representations of S_n .

The proof will use the following exercise.

Exercise 6.4.7. Let T be any $[n]$ -filling of λ , and let \mathfrak{U} be a tabloid of shape μ , where $\mu >_{\text{lex}} \lambda$. Prove that there must be entries $a \neq b$ in the same column of T and the same row of \mathfrak{U} .

Proof. By the preceding exercise and Lemma 6.4.3 we see that $b_T x = 0$ for all $x \in M^\mu$. If $\phi \in \text{Hom}(S^\lambda, M^\mu)^{S_n}$, then we have

$$\phi(v_T) = \phi\left(b_T \frac{v_T}{|C(T)|}\right) = b_T \phi\left(\frac{v_T}{|C(T)|}\right) = 0.$$

Since the vectors v_T span S^λ , ϕ is the zero map. □

6.5 Characters of Specht modules

Let χ^λ denote the character of S^λ .

Theorem 6.5.1. *The character of S^λ is given by $\chi^\lambda(\mu) = \langle s_\lambda, p_\mu \rangle$.*

Proof. Define a linear map Ψ from the space of class functions to the ring of symmetric functions by $\Psi(\xi^\lambda) = h_\lambda$. Since the characters ξ^λ are a basis for the space of class functions, every character χ_V must satisfy $\chi_V(\mu) = \langle \Psi(\chi_V), p_\mu \rangle$.

Now Ψ is an isometry, since

$$\begin{aligned} \langle \xi^\lambda, \xi^\nu \rangle &= \sum_{\mu \vdash n} \frac{1}{z(\mu)} \xi^\lambda(\mu) \xi^\nu(\mu) \\ &= \sum_{\mu \vdash n} \frac{1}{z(\mu)} \langle h_\lambda, p_\mu \rangle \langle h_\nu, p_\mu \rangle \\ &= \langle h_\lambda, h_\nu \rangle. \end{aligned}$$

Since the character ξ^μ is the sum of the characters of the irreducible subrepresentations of M^λ ,

$$\xi^\mu = \sum_{\lambda \leq_{\text{lex}} \mu} K_{\lambda\mu} \chi^\lambda,$$

for some integers $K_{\lambda\mu}$, where $K_{\lambda\lambda} > 0$. In other words, the change of basis matrix from $\{\chi^\lambda \mid \lambda \vdash n\}$ to $\{\xi^\lambda \mid \lambda \vdash n\}$ is upper triangular with positive diagonal entries, if the partitions of n are listed in lexicographically decreasing order. Moreover, we know that $\{\chi^\lambda\}$ is orthonormal basis for vector space of class functions. From linear algebra, we know that this is actually the unique basis with these properties: it is the basis obtained by applying the Gram-Schmidt orthogonalization process to $\{\xi^\mu\}$ (with partitions in lexicographically decreasing order).

Thus $\Psi(\chi^\lambda)$ is the result of applying Gram-Schmidt to the basis h_λ ; by the same argument as above, this is s_λ . and so $\chi^\lambda(\mu) = \langle \Psi(\chi^\lambda), p_\mu \rangle = \langle s_\lambda, p_\mu \rangle$. □

This argument also shows that the numbers $K_{\lambda\mu}$ above are the Kostka numbers. We can also conclude that all characters of S_n -representations are integer valued functions, since s_λ can be expressed as an integer linear combination of h_λ . Another easy consequence is the dimension of S^λ .

Corollary 6.5.2. $\dim S^\lambda = f^\lambda$.

Proof. The dimension of S^λ is $\chi^\lambda(\text{id}) = \chi^\lambda(1^n) = \langle s_\lambda, p_{1^n} \rangle = \langle s_\lambda, h_{1^n} \rangle$, which is f^λ , as we have seen earlier. \square

Exercise 6.5.3. Prove that the set $\{v_T \mid T \in \text{SYT}(\lambda)\}$ is a basis for S^λ .

On a related note, we can give a quick representation-theoretic proof of the identity

$$n! = \sum_{\lambda \vdash n} (f^\lambda)^2. \quad (6.5.1)$$

The right hand side is $\sum_{\lambda \vdash n} \langle s_\lambda, p_{1^n} \rangle^2 = \langle p_{1^n}, p_{1^n} \rangle = z(1^n)$, which equals $n!$.

Theorem 6.5.4. *The group algebra of the symmetric group $\mathbb{C}[S_n]$ is isomorphic as an algebra to the direct sum of matrix algebras*

$$\bigoplus_{\lambda \vdash n} \text{End}(S^\lambda),$$

where the map from $\mathbb{C}[S_n]$ to $\text{End}(S^\lambda)$ is just ρ_{S^λ} .

This gives another representation-theoretic interpretation of the identity (6.5.1). The left hand side is $\dim \mathbb{C}[S(n)] = n!$, whereas $\dim \text{End}(S^\lambda) = (f^\lambda)^2$, hence $\dim \bigoplus_{\lambda \vdash n} \text{End}(S^\lambda)$ is given by the right hand side.

Proof. Since the dimensions of both algebras are the same, it is enough to show that the map between them is injective. This is equivalent to proving that $\rho_{S^\lambda}(a) \neq 0$ for some $\lambda \vdash n$, if $a \in \mathbb{C}[S_n]$ is not the zero element.

We begin by viewing $\mathbb{C}[S_n]$ as a representation of S_n , where $\rho_{\mathbb{C}[S_n]}(\pi)$ is the left multiplication operator, defined by $\rho_{\mathbb{C}[S_n]}(\pi)x = \pi x$. This is called the left-regular representation S_n ; it is isomorphic to the representation associated to the species \mathcal{L} of linear orders. We therefore obtain a linear map $\rho_{\mathbb{C}[S_n]} : \mathbb{C}[S_n] \rightarrow \text{End}(\mathbb{C}[S_n])$. It is clear that $\rho_{\mathbb{C}[S_n]}(a) \neq 0$ if $a \neq 0$, since $a = \rho_{\mathbb{C}[S_n]}(a)(\text{id})$.

But now, if we decompose $\mathbb{C}[S_n]$ into irreducible subrepresentations, then for each $0 \neq a \in \mathbb{C}[S_n]$, $\rho_{\mathbb{C}[S_n]}(a)$ decomposes as a block diagonal operator, where each diagonal block is similar to $\rho_{S^\lambda}(a)$ for some $\lambda \vdash n$. Since $\rho_{\mathbb{C}[S_n]}(a) \neq 0$, we can find a non-zero block, which gives us $\rho_{S^\lambda}(a) \neq 0$. \square

Still another interpretation of (6.5.1) is the fact that the left-regular representation of S_n decomposes into f^λ copies of each irrep S^λ . Both this, and Theorem 6.5.4 generalize to any finite group.

It is possible to compute the character χ^λ recursively. We say that a skew partition λ/μ is a connected ribbon, if there is unique path between any two boxes with horizontal and vertical steps. In particular λ/μ must come in ‘‘one piece’’ and cannot contain a 2×2 square of boxes. We define $\delta(\lambda/\mu)$ to be the number of nonempty rows in λ/μ .

Theorem 6.5.5. *Let μ' be obtained by removing a part of size k from μ . Then*

$$\chi^\lambda(\mu) = \sum (-1)^{\delta(\lambda/\lambda')-1} \chi^{\lambda'}(\mu')$$

where the sum is taken over all λ' such that λ/λ' is a connected ribbon of size k .

Exercise 6.5.6. Prove Theorem 6.5.5. Here's one possible approach. First prove it for $\lambda' = \epsilon$, to express p_k in terms of Schur functions. Then use the Littlewood-Richardson rule to prove that $s_{\lambda'} p_k = \sum (-1)^{\delta(\lambda/\lambda')-1} s_{\lambda}$, where the sum is taken over all λ such that λ/λ' is a connected ribbon. Deduce the theorem from this formula.

6.6 The centre of the group algebra

One place where characters of S_n arise is in the study of the centre of $\mathbb{C}[S_n]$. Let

$$C_{\mu} = \sum_{\pi \in \mathcal{C}_{\mu}} \pi \in \mathbb{C}[S_n]$$

be the sum of all permutations of cycle type μ . Then for any $\pi \in S_n$, we have $\pi C_{\mu} \pi^{-1} = C_{\mu}$, since $\pi \sigma \pi^{-1}$ has cycle type μ iff σ has cycle type μ . Thus C_{μ} is in the centre of $\mathbb{C}[S_n]$.

Proposition 6.6.1. *The centre of $\mathbb{C}[S_n]$ has basis $\{C_{\mu} \mid \mu \vdash n\}$.*

Proof. It is clear that $\{C_{\mu} \mid \mu \vdash n\}$ is linearly independent. To see that these are spanning, we compute the dimension of the centre, using the isomorphism $\mathbb{C}[S_n] \simeq \bigoplus_{\lambda \vdash n} \text{End}(S_{\lambda})$. On the right hand side, the centre is given by all block diagonal operators where each diagonal block is a multiple of the identity. This shows that the dimension of the centre is the number of partitions of n . \square

Here is another approach to showing that $\{C_{\mu}\}$ spans the centre of $\mathbb{C}[S_n]$. View $\mathbb{C}[S_n]$ as a representation of S_n , where $\rho_{\mathbb{C}[S_n]}(\pi)x = \pi x \pi^{-1}$. This is isomorphic to the representation associated to the species \mathcal{S} of permutations; it is not isomorphic to the left-regular representation! (This is also a case where we strenuously avoid writing πx for this action, since it could be easily confused with multiplication in the group algebra. If a shorthand is needed, one might write $\pi \cdot x = \pi x \pi^{-1}$.) Then the center of $\mathbb{C}[S_n]$ is $\mathbb{C}[S_n]^{S_n}$. The projection operator $P_{\mathbb{C}[S_n]}$ of Proposition 6.2.4 maps $\pi \in S_n$ to $\frac{1}{z(\mu)} C_{\mu}$ if $\pi \in \mathcal{C}_{\mu}$, which tells us that these are spanning.

Proposition 6.6.2. *If $x \in S^{\lambda}$, then $C_{\mu} x = \alpha_{\mu\lambda} x$, where*

$$\alpha_{\mu\lambda} = \frac{n!}{z(\mu) f^{\lambda}} \chi^{\lambda}(\mu). \quad (6.6.1)$$

Proof. Since C_{μ} is in the the centre of $\mathbb{C}[S_n]$, $\rho_{S^{\lambda}}(C_{\mu})$ is in the centre of $\text{End}(S_{\lambda})$ and so it must equal $\alpha_{\mu\lambda} I_{S_{\lambda}}$ for some scalar $\alpha_{\mu\lambda}$. To compute $\alpha_{\mu\lambda}$, we use the fact that each of the $\frac{n!}{z(\lambda)}$ terms in the sum defining C_{μ} has the same character $\chi^{\lambda}(\mu)$. This gives us

$$\alpha_{\mu\nu} \dim S^{\lambda} = \text{tr}(\rho_{S^{\lambda}}(C_{\mu})) = \frac{n!}{z(\mu)} \chi^{\lambda}(\mu).$$

from which the result follows. \square

Since the characters χ^{λ} are linearly independent, $(\alpha_{\mu\lambda})_{\mu, \lambda \vdash n}$ defines an invertible matrix. Let $(\beta_{\theta\nu})_{\theta, \nu \vdash n}$ denote the inverse matrix.

Proposition 6.6.3.

$$\beta_{\theta\nu} = \frac{f^\theta}{n!} \chi^\theta(\nu) \quad (6.6.2)$$

Proof. We just need to check that $\sum_{\theta \vdash n} \alpha_{\mu\theta} \beta_{\theta\nu} = \delta_{\mu,\nu}$; and indeed

$$\begin{aligned} \sum_{\theta \vdash n} \alpha_{\mu\theta} \beta_{\theta\nu} &= \sum_{\theta \vdash n} \frac{1}{z(\mu)} \chi^\theta(\mu) \chi^\theta(\nu) \\ &= \frac{1}{z(\mu)} \sum_{\theta \vdash n} \langle s_\theta, p_\mu \rangle \langle s_\theta, p_\nu \rangle \\ &= \frac{1}{z(\mu)} \langle p_\mu, p_\nu \rangle = \delta_{\mu,\nu}. \quad \square \end{aligned}$$

Since $(\beta_{\theta\nu})$ is an invertible matrix, we can define a new basis for the centre of $\mathbb{C}[S_n]$:

$$F_\theta = \sum_{\nu \vdash n} \beta_{\theta\nu} C_\nu. \quad (6.6.3)$$

Proposition 6.6.4. F_θ are orthogonal idempotents, which means that

$$F_\theta F_{\theta'} = \delta_{\theta\theta'} F_\theta.$$

Proof. For $x \in S^\lambda$ we have

$$F_\theta x = \sum_{\nu \vdash n} \beta_{\theta\nu} \alpha_{\nu\lambda} x = \delta_{\lambda,\theta} x$$

and moreover F_θ is the unique element in $\mathbb{C}[S_n]$ with this property. The result follows by considering how the both sides act on $x \in S^\lambda$ for all $\lambda \vdash n$. \square

Inverting (6.6.3) we obtain

$$C_\mu = \sum_{\theta \vdash n} \alpha_{\mu\theta} F_\theta,$$

which allows us to compute products of C_μ , in terms of characters.

Theorem 6.6.5.

$$C_\mu C_\nu = \frac{n!}{z(\mu)z(\nu)} \sum_{\lambda \vdash n} \left(\sum_{\theta \vdash n} \frac{\chi^\theta(\mu) \chi^\theta(\nu) \chi^\theta(\lambda)}{f^\theta} \right) C_\lambda$$

Proof.

$$C_\mu C_\nu = \sum_{\theta \vdash n} \alpha_{\mu\theta} \alpha_{\nu\theta} F_\theta = \sum_{\lambda \vdash n} \sum_{\theta \vdash n} \alpha_{\mu\theta} \alpha_{\nu\theta} \beta_{\theta\lambda} C_\lambda$$

The result now follows by using the formulae (6.6.1) and (6.6.2) to evaluate $\alpha_{\mu\theta}$, $\alpha_{\nu\theta}$, $\beta_{\theta\lambda}$. \square

One reason to consider this formula is that it tells us how many pairs of permutations $(\sigma, \tau) \in \mathcal{C}_\mu \times \mathcal{C}_\nu$ have their product in \mathcal{C}_λ . Equivalently, this is the number of triples $(\pi, \sigma, \tau) \in \mathcal{C}_\lambda \times \mathcal{C}_\mu \times \mathcal{C}_\nu$ whose product is the identity element. Let

$$\begin{aligned} a_{\lambda,\mu,\nu} &= \#\{(\sigma, \tau) \in \mathcal{C}_\mu \times \mathcal{C}_\nu \mid \sigma\tau \in \mathcal{C}_\lambda\} \\ &= \#\{(\pi, \sigma, \tau) \in \mathcal{C}_\lambda \times \mathcal{C}_\mu \times \mathcal{C}_\nu \mid \pi\sigma\tau = \text{id}\}. \end{aligned}$$

Corollary 6.6.6.

$$a_{\lambda,\mu,\nu} = \frac{(n!)^2}{z(\lambda)z(\mu)z(\nu)} \sum_{\theta \vdash n} \frac{\chi^\theta(\lambda)\chi^\theta(\mu)\chi^\theta(\nu)}{f^\theta} \quad (6.6.4)$$

Note that the formula (6.6.4) is symmetrical in λ , μ and ν , as we would expect from the second definition of $a_{\lambda,\mu,\nu}$.

Proof. For each $\pi \in \mathcal{C}_\lambda$, the number of pairs $(\sigma, \tau) \in \mathcal{C}_\mu \times \mathcal{C}_\nu$ such that $\sigma\tau = \pi$ is given by

$$[\pi]\mathcal{C}_\mu\mathcal{C}_\nu = [\mathcal{C}_\lambda]\mathcal{C}_\mu\mathcal{C}_\nu = \frac{n!}{z(\mu)z(\nu)} \sum_{\theta \vdash n} \frac{\chi^\theta(\mu)\chi^\theta(\nu)\chi^\theta(\lambda)}{f^\theta}$$

Summing over all $\frac{n!}{z(\lambda)}$ possible values for π , gives the result. \square

It is convenient to put the numbers $a_{\lambda,\mu,\nu}$ into an exponential generating function. Consider the species of \mathcal{Z} of all triples of permutations (π, σ, τ) such that $\pi\sigma\tau = \text{id}$. We define a weight function by

$$\text{WT}(\pi, \sigma, \tau) = p_\lambda(w)p_\mu(x)p_\nu(y), \quad (6.6.5)$$

if $(\pi, \sigma, \tau) \in \mathcal{C}_\lambda \times \mathcal{C}_\mu \times \mathcal{C}_\nu$, where $w_1, w_2, \dots, x_1, x_2, \dots, y_1, y_2, \dots$ are three different sets of variables. This produces the generating function

$$1 + \sum_{n \geq 1} \frac{t^n}{n!} \sum_{\lambda, \mu, \nu \vdash n} a_{\lambda, \mu, \nu} p_\lambda(w) p_\mu(x) p_\nu(y). \quad (6.6.6)$$

Notice that the variable t records redundant information: the size of the partitions λ, μ, ν are already encoded in p_λ, p_μ, p_ν . Hence, we put $t = 1$, to obtain

$$A(w, x, y) = 1 + \sum_{n \geq 1} \frac{1}{n!} \sum_{\lambda, \mu, \nu \vdash n} a_{\lambda, \mu, \nu} p_\lambda(w) p_\mu(x) p_\nu(y).$$

We recover (6.6.6) by considering $A(tw, x, y)$, where $tw = (tw_1, tw_2, \dots)$.

The coefficients $a_{\lambda,\mu,\nu}$ are then given by

$$a_{\lambda,\mu,\nu} = \frac{\langle A(w, x, y), p_\lambda(w)p_\mu(x)p_\nu(y) \rangle_{w,x,y}}{z(\lambda)z(\mu)z(\nu)},$$

where $\langle \cdot, \cdot \rangle_{w,x,y}$ is computed by taking the Hall inner product with respect to each of the three sets of variables: i.e. $\langle \cdot, f(w)g(x)h(y) \rangle_{w,x,y} = \langle \langle \cdot, f(w) \rangle_w, g(x) \rangle_x, h(y) \rangle_y$.

Theorem 6.6.7.

$$A(w, x, y) = \sum_{\theta} \frac{|\theta|!}{f^\theta} s_\theta(w) s_\theta(x) s_\theta(y)$$

where the sum is taken over all partitions θ .

Proof. From (6.6.4), we have

$$\begin{aligned}
A(w, x, y) &= 1 + \sum_{n \geq 1} \frac{1}{n!} \left(\frac{(n!)^2}{z(\mu)z(\nu)z(\lambda)} \sum_{\theta \vdash n} \frac{\chi^\theta(\lambda)\chi^\theta(\mu)\chi^\theta(\nu)}{f^\theta} \right) p_\lambda(w)p_\mu(x)p_\nu(y) \\
&= 1 + \sum_{n \geq 1} \sum_{\theta \vdash n} \frac{n!}{f^\theta} \left(\sum_{\lambda \vdash n} \frac{\chi^\theta(\lambda)p_\lambda(w)}{z(\lambda)} \right) \left(\sum_{\mu \vdash n} \frac{\chi^\theta(\mu)p_\mu(x)}{z(\mu)} \right) \left(\sum_{\nu \vdash n} \frac{\chi^\theta(\nu)p_\nu(y)}{z(\mu)} \right) \\
&= 1 + \sum_{n \geq 1} \sum_{\theta \vdash n} \frac{n!}{f^\theta} s_\theta(w)s_\theta(x)s_\theta(y) \\
&= \sum_{\theta} \frac{|\theta|!}{f^\theta} s_\theta(w)s_\theta(x)s_\theta(y) \quad \square
\end{aligned}$$

This is one example of a situation where the use of power sum symmetric functions in a weight function leads to a nice generating function. Another example is the following exercise.

Exercise 6.6.8. Consider the weight function on the species \mathcal{S} of permutations, given by $\text{WT}(\pi) = p_\lambda$ if $\pi \in \mathcal{C}_\lambda$. Show that the mixed generating function

$$1 + \sum_{n \geq 1} \left(\sum_{\pi \in \mathcal{S}_n} \text{WT}(\pi) \right) \frac{t^n}{n!}$$

for \mathcal{S} is just $H(t) = \sum_{n \geq 0} h_n t^n$.

The species \mathcal{Z} is not a connected species; each object $(\pi, \sigma, \tau) \in \mathcal{Z}_X$ can be decomposed into “connected components”. To describe what these components look like, recall that a permutation can be represented by a directed graph, which is a union of directed cycles (which are the connected components for the permutation). Consider the graph G obtained by superimposing the three directed graphs for π , σ , and τ . We’ll say that (π, σ, τ) is connected if G is a connected graph; if it is not connected, the components of (π, σ, τ) correspond to the components the graph G . If X_1, \dots, X_k are the vertex sets of each component of G , then each X_i is invariant under π, σ and τ . The connected component of (π, σ, τ) corresponding to X_i is the connected triple of permutations $(\pi|_{X_i}, \sigma|_{X_i}, \tau|_{X_i}) \in \mathcal{Z}_{X_i}$.

Consider the species $\tilde{\mathcal{Z}}$ consisting of connected triples (τ, π, σ) such that $\tau\pi\sigma = 1$, endowed weight function (6.6.5). Since every \mathcal{Z} structure decomposes uniquely into connected components, we have a weighted natural equivalence $\mathcal{Z} \simeq \mathcal{E}[\tilde{\mathcal{Z}}]$, and hence the generating function for $\tilde{\mathcal{Z}}$ is

$$\tilde{A}(w, x, y) = \log A(w, x, y) = \log \left(\sum_{\theta} \frac{|\theta|!}{f^\theta} s_\theta(w)s_\theta(x)s_\theta(z) \right).$$

The number $\tilde{a}_{\lambda, \mu, \nu}$ of connected triples $(\tau, \pi, \sigma) \in \mathcal{C}_\lambda \times \mathcal{C}_\mu \times \mathcal{C}_\nu$ such that $\tau\pi\sigma = \text{id}$ is then given by

$$\tilde{a}_{\lambda, \mu, \nu} = \frac{\langle \tilde{A}(w, x, y), p_\lambda(w)p_\mu(x)p_\nu(y) \rangle_{w, x, y}}{z(\lambda)z(\mu)z(\nu)}. \quad (6.6.7)$$

6.7 Enumeration of maps

As an application, we consider the problem of enumerating *labelled maps* in a closed connected oriented surface. A map is a graph embedded in the surface so that all faces are two-cells (homeomorphic to a disc). We allow the possibility of loops and multiple edges. If a map has n edges, a labelling assigns labels $1, 2, \dots, 2n$ to each of the (two) ends of the edges.




Figure 6.7.1: A labelled map on a torus.

Example 6.7.1. A labelled map in the torus is drawn in Figure 6.7.1. The usual parallel line convention is used for the rectangular boundary, instructing us to identify top and bottom, in the same direction, and left and right, again in the same direction. The map has 4 vertices, 7 edges, and 3 faces (two of the faces are triangular, drawn in the middle of the rectangle, the remaining face consists of the four corner regions in the drawing). The 14 ends of the edges are labelled $1, \dots, 14$. As a consistency check, note that Euler's formula gives

$$V - E + F = 4 - 7 + 3 = 0 = 2 - 2g,$$

and indeed we have genus $g = 1$ for the torus.

Now let π be the permutation whose disjoint cycles list the labels encountered when moving around each vertex in counterclockwise order. Let σ be the permutation whose disjoint cycles (all of length 2) list the pairs of labels for the edges. Let τ be the permutation whose disjoint cycles list the labels encountered at the head of each edge when moving around each face in counterclockwise order. For example, in Figure 6.7.1, we obtain

$$\begin{aligned}\pi &= (1\ 4\ 5\ 11)(2\ 13\ 7)(3\ 12\ 8\ 9)(6\ 4\ 10) \\ \sigma &= (1\ 13)(2\ 12)(3\ 4)(5\ 14)(6\ 7)(8\ 11)(9\ 10) \\ \tau &= (1\ 8\ 2\ 6\ 9\ 11\ 14\ 7)(3\ 10\ 5)(4\ 13\ 12)\end{aligned}$$

Note that $\pi\sigma\tau = \text{id}$. This is because moving forward around a face, then back along an edge, then back around a vertex always returns to where we started. Note also that $\pi \in \mathcal{C}_{4433}$ and

$\tau \in \mathcal{C}_{833}$ precisely because the vertex degrees of the map are specified by the parts of the partition 4433, and the face degrees of the map are specified by the parts of the partition 833. Since there are seven edges and each edge has two ends, we have $\sigma \in \mathcal{C}_{27}$.

Given any labelled map with n edges, we can construct permutations τ, π, σ with $\tau\pi\sigma = \text{id}$, as in Example 6.7.1. The cycle types of π and τ give the degree sequences of the vertices and the faces in the map, respectively. The permutations π, σ and τ are enough to uniquely reconstruct the labelled drawing of the map. To see this, note that the permutations σ and τ together determine the labelling on each face; and π determines how those faces must fit together at each vertex.

However, this is not a bijective correspondence. The problem is that for some (π, σ, τ) we obtain a disconnected map when we attempt the reconstruction process. This will occur if and only if (π, σ, τ) is a disconnected triple of permutations. Hence what we have is a bijection between labelled maps and connected (π, σ, τ) such that $\pi\sigma\tau = \text{id}$, where $\sigma \in \mathcal{C}_{2n}$, and the cycle type of τ and π are the face- and vertex-degree sequences of the map.

Putting it all together, the number of labelled maps with vertex-degree sequence λ and face-degree sequence μ is given by $\tilde{a}_{\lambda, 2n, \mu}$, which can be computed by (6.6.7).

Chapter 7

Representations of $\mathfrak{sl}_2(\mathbb{C})$

Lie algebras are another algebraic object whose representation theory has many connections to combinatorics. In this chapter, we'll look at some of the combinatorial aspects and applications of the representation theory of the Lie algebra $\mathfrak{sl}_2(\mathbb{C})$.

7.1 What is an $\mathfrak{sl}_2(\mathbb{C})$ -representation?

A (finite dimensional) representation of $\mathfrak{sl}_2(\mathbb{C})$ is a finite dimensional complex vector space V together with three linear operators E, F, H on V satisfying

$$EF - FE = 2H, \quad HE - EH = 2E, \quad HF - FH = -2F. \quad (7.1.1)$$

A subspace $W \subset V$ is an $\mathfrak{sl}_2(\mathbb{C})$ -subrepresentation if W is an invariant subspace for E, F and H .

To understand where this definition is coming from, we need to explain what a Lie algebra is. This is a rather different kind of algebraic structure from the types of algebraic structures we have considered so far. A *Lie algebra* is a vector space \mathfrak{g} with a bilinear operation $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$, satisfying two conditions:

$$\begin{aligned} [x, y] &= -[y, x] && \text{for all } x, y \in \mathfrak{g}, \\ [x, [y, z]] &= [[x, y], z] + [y, [x, z]] && \text{for all } x, y, z \in \mathfrak{g}. \end{aligned}$$

It is best not to think of this operation as a “product”, since it is not an associative operation, and so it does not produce a ring structure on \mathfrak{g} . Instead it is called a *Lie bracket*.

The Lie algebra $\mathfrak{sl}_2(\mathbb{C})$ is defined to be the three dimensional vector space of 2×2 complex matrices of the form $\begin{pmatrix} a & b \\ c & -a \end{pmatrix}$, where the Lie bracket is defined by $[x, y] = xy - yx$. Note that in general, neither the matrix product xy nor yx is in $\mathfrak{sl}_2(\mathbb{C})$, but their difference always is. A basis for $\mathfrak{sl}_2(\mathbb{C})$ is

$$e = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \quad f = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} \quad h = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

and these 2×2 matrices satisfy the relations

$$ef - fe = 2h, \quad he - eh = 2e, \quad hf - fh = -2f.$$

7.2 Main example

Our main example of an $\mathfrak{sl}_2(\mathbb{C})$ -representation will be the following:

Example 7.2.1. Let $n \geq 0$ be an integer, and consider the vector space \mathbb{C}^{2^n} , which has a basis

$$\{x_\sigma \mid \sigma \in \{1, 2\}^n\}.$$

With think of $\sigma = \sigma_1 \dots \sigma_n$ as a string of 1's and 2's. To define a linear operator, it is enough to define it on basis elements. For this, we think of $\{1, 2\}^n$ as a poset: $\sigma \leq \tau$ if $\sigma_i \leq \tau_i$ for all i . Then

$$\begin{aligned} E(x_\sigma) &= \sum_{\tau \prec \sigma} x_\tau \\ F(x_\sigma) &= \sum_{\tau \succ \sigma} x_\tau \\ H(x_\sigma) &= \left(\#1\text{'s in } \sigma - \#2\text{'s in } \sigma \right) x_\sigma. \end{aligned}$$

For example, if $n = 2$ then

$$\begin{array}{lll} E(x_{11}) = 0 & F(x_{11}) = x_{12} + x_{21} & H(x_{11}) = 2x_{11} \\ E(x_{12}) = F(x_{21}) = x_{11} & F(x_{12}) = E(x_{21}) = x_{22} & H(x_{12}) = F(x_{21}) = 0 \\ E(x_{22}) = x_{12} + x_{21} & F(x_{22}) = 0 & H(x_{22}) = -2x_{22} \end{array}$$

It is easy to check that $EH(x_\sigma) - HE(x_\sigma) = 2E(x_\sigma)$ and $FH(x_\sigma) - HF(x_\sigma) = -2F(x_\sigma)$. To determine $EF(x_\sigma) - FE(x_\sigma)$, we need to count the number of sequences

$$\sigma, \tau, v \quad \text{where } \sigma \prec \tau \succ v$$

and subtract the number of sequences

$$\sigma, \tau', v \quad \text{where } \sigma \succ \tau' \prec v.$$

If $\sigma \neq v$ there are the same number in both cases, either 0 or 1, and so the difference is zero. If $\sigma = v$, the number of the first kind is the number of 1's in σ ; while the number of the second kind is the number of 2's in σ . Hence we see that $(EF - FE)(x_\sigma) = H(x_\sigma)$. Therefore \mathbb{C}^{2^n} is a representation of $\mathfrak{sl}_2(\mathbb{C})$

We can also define a representation of the symmetric group S_n on \mathbb{C}^{2^n} by putting

$$\pi x_{\sigma_1 \dots \sigma_n} = x_{\sigma_{\pi(1)} \dots \sigma_{\pi(n)}},$$

for $\pi \in S_n$.

Proposition 7.2.2. *If G is any subgroup of S_n , then the G -invariant vectors in \mathbb{C}^{2^n}*

$$W = \{x \in \mathbb{C}^{2^n} \mid gx = x \text{ for all } g \in G\}$$

are an $\mathfrak{sl}_2(\mathbb{C})$ -subrepresentation \mathbb{C}^{2^n} .

Proof. It is not hard to see that for any $x \in \mathbb{C}^{2^n}$, $\pi \in S_n$, we have

$$\pi E(x) = E(\pi x) \quad \pi F(x) = F(\pi x) \quad \pi H(x) = H(\pi x).$$

For any $x \in W$ and any $g \in G$ we have $gE(x) = E(gx) = E(x)$ so $E(x) \in W$. Similarly $F(x), H(x) \in W$, so W is a $\mathfrak{sl}_2(\mathbb{C})$ -subrepresentation \mathbb{C}^{2^n} . \square

7.3 Palindromic and unimodal sequences

Theorem 7.3.1. *Given a representation V of $\mathfrak{sl}_2(\mathbb{C})$ the following are true:*

(i) *H is diagonalizable and its eigenvalues are integers. Moreover, there is a basis $\mathcal{B} = \{z_1, \dots, z_N\}$ for V of eigenvectors for H , such that for $0 \leq i \leq N$, Ez_{i+1} is a non-negative integer multiple of z_i and Fz_i is a non-negative integer multiple of z_{i+1} (with the convention that $z_0 = z_{N+1} = 0$).*

We write $V = \bigoplus_{k \in \mathbb{Z}} V_k$, where V_k is the k -eigenspace of H .

(ii) *E maps V_k to V_{k+2} ; this is injective for $k \leq -1$, and surjective for $k \geq -1$. Similarly, F maps V_k to V_{k-2} ; this map is injective for $k \geq 1$, and surjective for $k \leq 1$.*

(iii) *For $k \geq 0$, E^k maps V_{-k} bijectively to V_k , and F^k maps V_k bijectively to V_{-k} .*

(iv) *For all $x \in V$: if $Ex \neq 0$ then $FEx \neq 0$; if $Fx \neq 0$ then $EFx \neq 0$.*

We'll call a basis $\{z_1, \dots, z_N\}$ for V as described in (i) a *good basis* for V . Given a good basis for V , form a directed graph Γ with vertex set $1, \dots, N$ and an arrow from $i+1$ to i iff $E(z_{i+1})$ is a *non-zero* multiple of z_i . (Equivalently, by (iv), $F(z_i)$ is a non-zero multiple of z_{i+1} .) Each component of this graph is a directed path.

Corollary 7.3.2. *From Γ , we can uniquely determine the operators H , E and F (up to diagonal change of basis).*

Proof. Note that the eigenvalue γ_i of H associated to z_i must increase by 2 as we move along each path (by (ii)), and must be symmetrical in each component (by (iii)). Thus the graph uniquely determines H . Assume that $Ez_i = \alpha_i z_{i-1}$ and $Fz_i = \beta_i z_{i+1}$. Then from $(EF - FE)(z_i) = H(z_i)$, we have

$$\alpha_{i+1}\beta_i - \beta_{i-1}\alpha_i = \gamma_i,$$

which allows us to write β_i in terms of α_i and γ_i . Changing the constants α_i is a diagonal change of basis. \square

A finite sequence a_0, \dots, a_n is called *palindromic* if $a_k = a_{n-k}$ for all $k \leq n$. It is called *unimodal* if for some $m \leq n$, $a_0 \leq a_1 \leq \dots \leq a_{m-1} \leq a_m$ and $a_m \geq a_{m+1} \geq \dots \geq a_{n-1} \geq a_n$. For example, the sequence 1, 1, 3, 8, 8, 3, 1, 1 is unimodal and palindromic; the sequence 1, 3, 9, 8, 8, 4, 4, 0 unimodal but not palindromic.

Corollary 7.3.3. *If $d_k = \dim V_k$ then for all $k \in \mathbb{Z}$, the sequences*

$$d_{-2k}, d_{-2k+2}, \dots, d_{-2}, d_0, d_2, \dots, d_{2k-2}, d_{2k}$$

and

$$d_{-2k-1}, d_{-2k+1}, \dots, d_{-3}, d_{-1}, d_1, d_3, \dots, d_{2k-1}, d_{2k+1}$$

are unimodal and palindromic.

Proof. This follows immediately from (ii) and (iii) above. \square

The proof of Theorem 7.3.1 is a challenging exercise in Linear algebra, which we will not undertake here. Instead, we'll discuss some of applications and examples.

Example 7.3.4. Let $\mathcal{P}_{k,\ell}$ be the set of partitions with at most k parts, and largest part at most ℓ . Let p_m be the number of partitions in $\mathcal{P}_{k,\ell}$ of size m . We've seen that

$$\sum_{m \geq 0} p_m q^m = \binom{k + \ell}{k}_q.$$

We'll now prove that the sequence $p_0, p_1, \dots, p_{k\ell}$ is palindromic and unimodal. It is quite easy to find a bijection that proves that this sequence is palindromic; it is not so easy to prove that it is unimodal. We'll do this by constructing an appropriate $\mathfrak{sl}_2(\mathbb{C})$ representation.

Consider $\mathbb{C}^{2^{k\ell}}$, where we think of the basis elements x_σ as being indexed now by a $k \times \ell$ array σ of 1's and 2's. To each σ we can associate a partition $\lambda(\sigma) \in \mathcal{P}_{k,\ell}$ as follows: first count the number of 1's in each row of σ , to produce a list of k non-negative integers; then sort these integers into a partition. For all $\lambda \in \mathcal{P}_{k,\ell}$, let

$$w_\lambda = \sum_{\sigma: \lambda(\sigma) = \lambda} x_\sigma,$$

and let $W \subset \mathbb{C}^{2^{k\ell}}$ be the subspace spanned by these vectors. It is not hard to check that W is of the form given in Proposition 7.2.2, where G is the subgroup that permutes each row individually, as well as the set of rows. Hence W is an \mathfrak{sl}_2 -representation. Moreover, each vector w_λ is an eigenvector of H with eigenvalue $2|\lambda| - k\ell$. It follows that the dimensions of the H -eigenspaces are $\dim W_{2m - k\ell} = p_m$, and so the unimodality follows from Corollary 7.3.3.

7.4 Crystal operators revisited

Now we'll describe a good basis for \mathbb{C}^{2^n} , and the associated graph Γ . To start, we consider the subspace W of S_n -invariant vectors of \mathbb{C}^{2^n} . This is an $(n + 1)$ -dimensional spaces, spanned by vectors which we'll call $z_{1^k 2^{n-k}}$, where $1^k 2^{n-k}$ denotes the string of k ones followed by $n - k$ twos.

$$z_{1^k 2^{n-k}} = \sum_{\substack{\sigma \text{ has } k \text{ 1's} \\ \text{and } n-k \text{ 2's}}} x_\sigma. \tag{7.4.1}$$

These have the property that

$$\begin{aligned}
Ez_{1^k 2^{n-k}} &= (n-k)z_{1^{k+1} 2^{n-k-1}} \\
Fz_{1^k 2^{n-k}} &= k z_{1^{k-1} 2^{n-k+1}} \\
Hz_{1^k 2^{n-k}} &= (2k-n)z_{1^k 2^{n-k}} .
\end{aligned} \tag{7.4.2}$$

In other words, this is a good basis for the subrepresentation W .

We extend this to a basis of all of \mathbb{C}^{2^n} by defining a vector z_σ for every $\sigma \in \{1, 2\}^n$. To do this, we define a linear maps

$$\varphi_k : \mathbb{C}^{2^m} \rightarrow \mathbb{C}^{2^{m+2}} \quad \text{for } 0 \leq k \leq m$$

by

$$\varphi_k(x_\sigma) = x_{\sigma_1 \dots \sigma_k 21 \sigma_{k+1} \dots \sigma_m} - x_{\sigma_1 \dots \sigma_k 12 \sigma_{k+1} \dots \sigma_m} .$$

For example $\varphi_1(x_{1211}) = x_{\underline{12}1211} - x_{1\underline{12}211}$. We now define

$$z_{\sigma_1 \dots \sigma_k 21 \sigma_{k+1} \sigma_{n-2}} = \phi_k(z_{\sigma_1 \dots \sigma_{n-2}}) .$$

This together with (7.4.1) defines z_σ for all σ .

Example 7.4.1. For example to compute z_{2112} , we note that $2112 = \underline{21}12$ is obtained by inserting a copy 21 into the string 12 at the 0th position. Therefore we have

$$\begin{aligned}
z_{2112} &= \varphi_0(z_{12}) \\
&= \varphi_0(x_{12} + x_{21}) \\
&= \varphi_0(x_{12}) + \varphi_0(x_{21}) \\
&= (x_{\underline{21}12} - x_{\underline{12}12}) + (x_{\underline{21}21} - x_{\underline{12}21})
\end{aligned}$$

To compute z_{2211} , we note that $2211 = \underline{22}11$, and so $z_{2211} = \varphi_1(z_{21})$. But 21 is obtained by inserting 21 into the empty string ϵ , so $z_{21} = \varphi_0(z_\epsilon) = \varphi_0(x_\epsilon) = x_{21} - x_{12}$. Hence

$$\begin{aligned}
z_{2211} &= \varphi_1(x_{21} - x_{12}) \\
&= \varphi_1(x_{21}) - \varphi_1(x_{12}) \\
&= (x_{\underline{22}11} - x_{\underline{21}21}) - (x_{\underline{12}12} - x_{\underline{11}22})
\end{aligned}$$

There may be more than one way to recursively construct σ by inserting 21 into a string of the form $1^k 2^\ell$, hence it is important to note the following.

Proposition 7.4.2. *For all $\sigma \in \{1, 2\}^n$, z_σ is well defined.*

Proof. Since φ_k is linear, it is enough to check that if $k_1 \leq k_2$ then

$$\varphi_{k_1}(\varphi_{k_2}(x_\sigma)) = \varphi_{k_2+2}(\varphi_{k_1}(x_\sigma))$$

which shows that changing the order of insertion doesn't affect the answer. \square

Proposition 7.4.3. For all $x \in \mathbb{C}^{2^m}$, $k \leq m$, we have

$$\varphi_k(E(x)) = E(\varphi_k(x)), \quad \varphi_k(F(x)) = F(\varphi_k(x)), \quad \varphi_k(H(x)) = H(\varphi_k(x)). \quad (7.4.3)$$

Proof. Again, since φ_k and E, F, H are linear, it is enough to check this for $x = x_\sigma$, which is straightforward. \square

Theorem 7.4.4. The basis $\{z_\sigma \mid \sigma \in \{1, 2\}^n\}$ is a good basis for \mathbb{C}^{2^n} , in some order.

Proof. To see that this is a basis, observe that

$$z_\sigma = x_\sigma + \sum_{\tau >_{\text{lex}} \sigma} C_{\tau\sigma} x_\tau$$

for some $C_{\tau\sigma}$ and therefore we have an upper triangular change of basis. The fact that it is a good basis (in some order) follows the fact that $\{z_{1^k 2^{n-k}}\}$ is a good basis for W , using (7.4.3). \square

We can also use (7.4.3) to determine the graph structure Γ . Suppose $E(z_\sigma) = \alpha z_\tau$. To determine the right hand side, we successively remove all 21's from σ to obtain σ' ; then (7.4.2) to compute $E(z_{\sigma'}) = \alpha z_{\tau'}$. Then reinsert the 21's into τ' to obtain τ .

Equivalently this can be described as follows. Scanning σ from right to left, cross out the 2's and 1's in pairs: whenever we see a 2, look for the last read uncrossed 1. If it exists, cross out the pair. Now change the leftmost uncrossed 2 to a 1, to produce τ . If this is not defined, $E(z_\sigma) = 0$.

This should look familiar. We have just shown:

Theorem 7.4.5. The arrows in Γ are given by $\sigma \rightarrow E_1(\sigma)$, where E_1 is the crystal raising operator.

Here we're viewing σ as a filling of a Ferrers diagram with a single row of length n . Although we defined the crystal operators on tableaux, the algorithm makes perfect sense for any filling.

This has an interesting interpretation, when we apply the the RSK correspondence to σ . Here we think of σ as a set of ordered pairs $\{(1, \sigma_1), \dots, (n, \sigma_n)\}$. The RSK correspondence produces a pair (P_σ, Q_σ) of tableaux of the same shape λ , where P_σ has entries ≤ 2 , and $Q_\sigma \in \text{SYT}(\lambda)$. Note that the restriction on the entries of P implies that λ has at most 2 rows.

Corollary 7.4.6. If $\sigma, \tau \in \{1, 2\}^n$, then

- (i) σ and τ are in the same component of Γ if and only if $Q_\sigma = Q_\tau$;
- (ii) if $\sigma \rightarrow \tau$ is an arrow in Γ , then $P_\tau = E_1(P_\sigma)$.

To prove this amounts to checking certain properties of the crystal operator E_1 , and we leave it as an exercise.

There are similar statements about the representation theory of the Lie algebra $\mathfrak{sl}_m(\mathbb{C})$, relating it to strings $\sigma \in \{1, \dots, m\}^n$, the crystal operators E_a , $a = 1, \dots, m-1$, and the the RSK-correspondence. However, these are much more difficult to state (the definition of a good basis is much more complicated in general), and well beyond the scope of this course.