Groups, Geometries and Representation Theory

David A. Craven

Spring Term, 2013

Contents

1 Combinatorics of partitions **1**

1.1 Previous knowledge about symmetric groups 1

1.2 Partitions 2

1.3 James's abacus 3

1.4 Tableaux 7

1.5 The Robinson–Schensted correspondence 11

1.6 Tabloids and polytabloids 13

1.7 Two orderings on partitions 15

1.8 An ordering on tabloids 16

2 Constructing the representations of the symmetric group **19**

2.1 M^λ and S^λ 19

2.2 The branching rule 22

2.3 The characteristic p case 23

2.4 Young's rule 28

3 The Murnaghan–Nakayama rule **32**

3.1 Skew tableaux 32

3.2 The Littlewood–Richardson rule 33

3.3 The Murnaghan–Nakayama rule 37

Chapter 1

Combinatorics of partitions

1.1 Previous knowledge about symmetric groups

A few facts about the symmetric group first that we need, that should be known from earlier courses. We will use cycle notation for the elements of the symmetric group S_n, i.e., $(1,2,5)(3,4)$. We multiply left to right, so that $(1,2)(1,3) = (1,2,3)$.

Writing a permutation $\sigma \in S_n$ as a product of disjoint cycles, the *cycle type* of σ, written $1^{a_1}2^{a_2}\dots n^{a_n}$, is the string where a_i is the number of cycles of length i in σ. Hence $(1,2,3)(4,5,6)(7,8)(9) \in S_9$ has cycle type $1^1 2^1 3^2$.

Recall that conjugation acts by acting on the cycle itself by the conjugating element, so that $(1,2,3)(4,5)$, acted on by $(2,5,3)$, becomes $(1,5,2)(4,3)$. This immediately yields the following result.

Proposition 1.1 Two elements of S_n are conjugate if and only if they have the same cycle type.

Along with conjugacy classes come centralizers. The centralizer of a given cycle $(1,2,3,\dots,m)$ is the cyclic group of order m generated by the element itself. If a group element $g = (1,2,3)(4,5)$ say is made up of disjoint cycles all of differing lengths (i.e., the exponents in the cycle type above are all either 0 or 1) then its centralizer is generated by these different cycles, and so is the direct product of the cyclic subgroups generated by them. However, the centralizer of $(1,2)(3,4)$ in S_4 also includes the element $(1,3)(2,4)$ which swaps the cycles. In general we get the following theorem.

Theorem 1.2 The order of the centralizer of a cycle of type $1^{a_1}2^{a_2}\dots n^{a_n}$ with $\sum i \cdot a_i = n$ is

$$\prod_{i=1}^{n} i^{a_i} \cdot a_i!.$$

1.2 Partitions

Since the cycle type of an element of S_n is given by $1^{a_1}2^{a_2}\dots n^{a_n}$ with $\sum i \cdot a_i = n$, we have that the conjugacy classes of S_n are in one-to-one correspondence with *partitions* of n.

Definition 1.3 A *partition* of a non-negative integer n is a sequence $\lambda = (\lambda_1, \dots, \lambda_r)$ of positive integers such that $\lambda_i \geqslant \lambda_{i+1}$ and $\sum_i \lambda_i = n$. We write $n = |\lambda|$ is the *length* of λ, and often $\lambda \vdash n$.

The empty partition is the only partition of 0, and the partitions of four are (4), $(3,1)$, $(2,2)$, $(2,1,1)$ and $(1,1,1,1)$. In general the number of partitions grow sub-exponentially but super-polynomially.

Partitions can be represented via *Young diagrams*. These are essentially pictures with rows of boxes, whose ith row has λ_i boxes in it. For example, the partition $(4,4,3,1)$ is given by the diagram

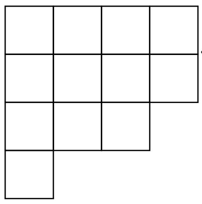.

Young diagrams are named after Alfred Young, and called Ferrers diagrams by some people, mostly combinatorialists. The combinatorics of Young diagrams is expansive, and we will only meet a small amount of it at the moment. The first object that is very easy to define using Young diagrams is the *conjugate partition*. Simply draw a line diagonally from the top-left box down and to the right, and then reflect the partition. One ends up with another partition of the same number, in the example above it is

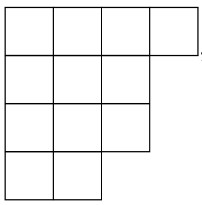,

and for example the partitions (n) and $(1,1,\dots,1)$ are swapped. This action induces a bijection on the set of all partitions of n, and this bijection has order 2. The conjugate partition has a variety of different notations in the literature, such as λ', λ^c, $\bar{\lambda}$, and we do not fix a particular notation here to maintain flexibility.

One of the most fundamental objects in the combinatorial theory of partitions are *hook lengths*. An algebraic definition is horrific, but a pictorial definition is fairly easy. Choose a box x in a Young diagram. The *hook* η_x consists of the box x, all boxes to the right of x, and all boxes below x. (If x has coordinates (i,j) then the hook is also written $\eta_{i,j}$.) The *arm* is those to the right, and the *leg* is all boxes below x. The *hook length*, *arm length* and *leg length* are the number of boxes in the hook, arm and leg respectively. It is clear that the hook length is the sum of the arm length, leg length, and 1. In the diagram above, adding hook lengths, we get

7	5	4	2
6	4	3	1
4	2	1	
1			

.

The hook lengths along the first column of the partition are rather uninspiringly called the *first-column hook lengths*. Notice that taking conjugate partitions swaps arms and legs, so leaves the hook length of (the reflection of) a box unchanged. We come to our first proposition about partitions.

Proposition 1.4 (i) If λ is a partition then the hook lengths along any row or column strictly increase towards the left and top respectively.

(ii) The act of taking first-column hook lengths induces a bijection between the collection of all partitions of all non-negative integers and the set of all finite subsets of the positive integers.

Proof: If a box x is further to the left than the box y but on the same row, then the arm length of x must be larger than that of y, and the leg length is at least as high, so the hook length of x is larger than that of y. Taking conjugates gives the same result for columns. This proves (i).

In order to prove (ii), we note that firstly by (i) the map is well defined. We show how to reconstruct the partitions λ from its first-column hook lengths $\{h_1, \dots, h_r\}$, and we order the h_i for simplicity so that $h_i > h_{i+1}$. The relationship between the h_i and the λ_i is given by

$$h_i = \lambda_i + r - i.$$

Since r is fixed, we can move between the λ_i and h_i, so the function of taking first-column hook lengths is indeed a bijection. $\square$

1.3 James's abacus

In this section we will introduce James's abacus and prove the uniqueness of t-cores. The abacus is a pictorial representation of the first-column hook lengths of a partition that can be manipulated to extract various combinatorial data about it, most notably its core.

Definition 1.5 Let λ be a partition and let t be a positive integer. A *t-hook* is a hook in λ of length t. The process of *removing* a t-hook consists of deleting the boxes that comprise a hook of length t, then moving the boxes that are currently below and to the right of the hook one box up and one box left, to construct another partition of $|\lambda| - t$. We write $\lambda \setminus \{\eta\}$ for the result of removing the hook η from the partition λ.

Algebraically, if $\lambda = (\lambda_1, \dots, \lambda_r)$ is a partition, and $x = (a, b)$ is a box in λ, then removing the hook η_x at x (which has leg length l) results in the partition $(\mu_1, \dots, \mu_r)$, where $\mu_i = \lambda_i$ for $i < a$, $\mu_i = \lambda_{i+1} - 1$ for $a \leqslant i < a + l$, $\mu_{a+l} = b - 1$, and $\mu_i = \lambda_i$ for $i > a + l$. If $i = 1$ then some of the μ_i might be zero, in which case these are removed.

The only part of the algebraic description that needs confirming is that $\mu_{a+l} = b-1$, and to see this we note that once the disconnected partition has been moved upward one row, there is nothing extra to add onto row $a+l$, so that row $a+l$ must end where the hook was taken away, i.e., $\mu_{a+l} = b - 1$.

Before we present James's abacus, we want to describe the effect of removing a hook from a partition on its set of first-column hook lengths. For this, the algebraic description comes in handy. The only problem is that the act of removing the zeros from the end of μ is difficult to describe, so we wish to leave them in; however, then we don't have a 'real' partition, and we need to expand our collection of partitions and first-column hook lengths to accommodate zeros.

Definition 1.6 A *β-set* (also called a set of β-numbers) is a finite collection of **non-negative** integers. We place an equivalence relation on the set of all β-sets by $X \sim Y$ if

$$Y = \{0\} \cup \{i + 1 \mid i \in X\},$$

and taking the reflexive, symmetric and transitive closure of this relation. In other words, $X \sim Y$ if one can get Y from X (or the other way round) by adding n to each element of X and appending $\{0, \dots, n-1\}$, for some $n \geqslant 0$.

As examples, the β-sets $\{0, 1, 4, 6, 8\}$ and $\{0, 3, 5, 7\}$ are equivalent, and $\{0, 1, 2\}$ is equivalent to the β-set $\emptyset$.

Inside each equivalence class of β-sets there exists one X for which $0 \notin X$, and X is the set of first-column hook lengths for a partition λ. (This is the empty partition when $X = \emptyset$.) Hence there is a bijection between partitions of non-negative integers and equivalence classes of β-sets, so to any β-set we can associate a partition, and to any partition we can associate *some* β-set.

A way of thinking of β-sets that are not first-column hook lengths for a partition is to take a partition λ, add finitely many zeroes to the end of λ, and then use for formula

$$x_i = \lambda_i + r - i$$

where λ has r parts, to produce a β-set of λ. It is easy to see that adding a zero to λ has the effect of increasing all the x_i by 1 and appending 0 to the end, so that this is indeed a β-set of λ.

We can now describe the process of removing a hook from λ in terms of β-sets as follows.

Proposition 1.7 Let $X = \{x_1, \dots, x_r\}$ be a β-set of a partition λ. Suppose that the hook at $x = (a, b)$ has length t. If Y is a β-set obtained by removing the hook at x, then (one choice of) Y is obtained from X by subtracting t from x_a and keeping all other x_i the same.

Proof: Let μ be the partition obtained from λ by removing the hook at x, and leave any zeroes at the end of μ. The algebraic definition of removing a t-hook is that $\lambda_i = \mu_i$ for $i < a$ and $i > a + l$, where l is the leg length of the hook at x, and $\mu_i = \lambda_{i+1} - 1$ for $a \leqslant i < a + l$, with $\mu_{a+l} = b - 1$. We will construct the β-set Y of μ that has the same number of elements as X.

Since the number of parts of λ and μ are the same, if $\lambda_i = \mu_i$ then $x_i = y_i$, so every element apart from those for $a \leqslant i \leqslant a + l$ are good. For $a \leqslant i < a + l$, we have

$$y_i = \mu_i + r - i = \lambda_{i+1} - 1 + r - i = \lambda_{i+1} + r - (i + 1) = x_{i+1},$$

and finally, since $\lambda_a = (b - 1) + t - l$ (the hook had length t, started at box b, and only the leg, with l boxes, is not in row a, which had λ_a boxes in total)

$$y_{a+l} = \mu_{a+l} + r - (a+l) = (b-1) + r - (a+l) = (\lambda_a - t + l) + r - a - l = (\lambda_a + r - a) - t = x_a - t.$$

Hence we see that Y can be obtained from X by removing t from x_a and leaving all other β-numbers the same. $\square$

We now introduce the abacus. Let X be a set of β-numbers. The abacus of X has t runners for some positive integer t, labelled runners $0, 1, \dots, t-1$, pointing downwards. Each runner consists of positions, one for each non-negative integer, and the value of position i at runner j is $i \cdot t + j$; so the value of positions on runner i are always congruent to i modulo t. Each position is occupied by a bead if the value of that position is in X, and a gap otherwise. Thus if $t = 3$ and $X = \{0, 1, 3, 5, 6, 8\}$, we have the following diagram.

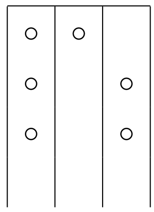

We can move from one β-set to an equivalent one by shifting all beads one to the right and adding a bead to the bottom, like this:

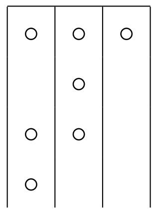

Proposition 1.7 proves that to obtain the β-set corresponding to removing a t-hook, we simply push a particular bead up one position on the abacus, like so:

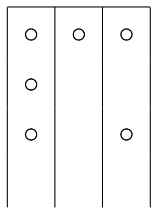

To see this in Young diagrams, we have the following hook removal:

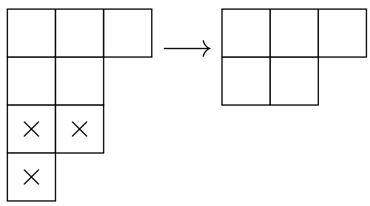

Notice that, if the β-set of the unmodified diagram is $\{1, 3, 4, 6\}$ then removing the 3-hook above yields $\{0, 1, 4, 6\}$, which is equivalent to $\{2, 4\}$, the first-column hook lengths of the partition on the right.

This process of removing a t-hook now makes well defined the following definition.

Definition 1.8 A *t-core* is a partition λ with no t-hooks. If λ is any partition, then the t-core of λ is the partition obtained by removing all t-hooks from λ.

A priori, it is not obvious that the order in which you remove the hooks does not matter. However, on the abacus, it is now clear that the t-core of a partition is obtained by simply shifting all beads as far up their runners as possible. The t-core of a partition is a fundamental object in the representation theory of the symmetric group, and we will see this importance (much) later in the course.

1.4 Tableaux

A *tableau* is a filling of the boxes in a Young diagram with the numbers 1 to n, as for example in the following.

$$\begin{array}{|c|c|c|}\hline 1 & 5 & 4 \\ \hline 3 & 2 & \multicolumn{1}{c}{} \\ \cline{1-2}\end{array}$$

We say that a tableau is *standard* if the numbers in the rows and columns are in increasing order. For example, the above tableau is not standard and the following tableau is standard.

$$\begin{array}{|c|c|c|}\hline 1 & 3 & 5 \\ \hline 2 & 4 & \multicolumn{1}{c}{} \\ \cline{1-2}\end{array}$$

The next lemma is clear, and no proof is needed.

Lemma 1.9 The number of Young tableaux for a given partition λ of size n is $n!$. There is at least one standard Young tableau with shape λ.

For almost all diagrams, i.e., not (n) and (1^n), there are at least two standard tableaux for that diagram.

$$\begin{array}{|c|c|c|}\hline 1 & 3 & 5 \\ \hline 2 & 4 & \multicolumn{1}{c}{} \\ \cline{1-2}\end{array} \qquad\qquad \begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline 4 & 5 & \multicolumn{1}{c}{} \\ \cline{1-2}\end{array}$$

We write f^λ for the number of standard tableau with shape λ.

Definition 1.10 A *removable node* in a Young diagram λ is one whose hook length is 1, i.e., a box which may be removed and the remaining daigram is that of a partition. An *addable node* is a removable node x in a Young diagram μ such that $\lambda = \mu \setminus \{x\}$. Write $\mathrm{Rem}(\lambda)$ for the set of all removable nodes of λ and $\mathrm{Add}(\lambda)$ for the set of addable nodes for λ.

Removable nodes are important since they allow us to compare partitions of n with partitions of $n-1$. Notice that if x is a box in a standard tableau, then all boxes in the arm or leg of the hook corresponding to x must have larger numbers than that of x, so whichever box is filled with n must be removable. Furthermore, removing this box, and its filling, results in a standard tableau for the Young diagram $\lambda \setminus \{x\}$.

What we see is that there is a bijection between the set of all standard tableaux with shape λ and the union, over all removable nodes x of λ, of the standard $\lambda \setminus \{x\}$-tableaux. In other words, we have

$$f^\lambda = \sum_{x \in \mathrm{Rem}(\lambda)} f^{\lambda \setminus \{x\}}.$$

Clearly we also have $f^{(n)} = f^{(1^n)} = 1$, as we mentioned before, so that any function g on the set of partitions that satisfies $g((1)) = 1$ and the recursion above, is a formula for f^λ. We will prove the remarkable identity

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

where $\lambda \vdash n$, (i,j) runs over all boxes in λ, and $h_{i,j}$ is the hook length of the box (i,j). The method of proof will be exactly as described, noting that clearly the right-hand side evaluates to 1 for the partition (1).

We first rewrite this in terms of the first-column hook lengths $h_{i,1}$.

Lemma 1.11 Let $\lambda = (\lambda_1, \ldots, \lambda_r)$ be a partition of n. We have

$$\prod_{(i,j)\in\lambda} h_{i,j} = \frac{\prod_{i=1}^{r}(h_{i,1})!}{\prod_{1\leqslant i<j\leqslant r}(h_{i,1} - h_{j,1})}.$$

Proof: For each $1 \leqslant i \leqslant r$ we have that there are no repetitions in the set

$$h_{i,1}, h_{i,2}, \ldots, h_{i,\lambda_i}, h_{i,1} - h_{i+1,1}, h_{i,1} - h_{i+2,1}, \ldots, h_{i,1} - h_{r,1},$$

i.e., the hook length of the box $x = (1,i)$, the hook lengths of the boxes in the arm of x, and the difference between the hook length of x and the hook lengths of the boxes in the leg of x. To see this, we first notice that removing the first $i-1$ rows from λ does not affect the numbers in the hook of the box, so we can assume that $i = 1$.

It is clear that no two hook lengths among x and the boxes in the arm of x are the same, and similarly no two hook lengths in the leg of x are the same. Furthermore, $h_{1,1} - h_{k,1}$ can never be equal to $h_{1,1}$, so that the only possible equality is $h_{1,j} = h_{1,1} - h_{k,1}$ for some $1 < j \leqslant \lambda_1$ and $1 < k \leqslant r$. Writing l for the leg length of $(1,j)$, we get that $h_{1,1} = \lambda_1 + r - 1$, $h_{k,1} = \lambda_k + r - k$, and $h_{1,j} = \lambda_1 - j + 1 + l$, so the equality becomes

$$(j-1) + (k-1) = l + \lambda_k.$$

If $\lambda_k \leqslant j-1$ then the hook at $(1,j)$ cannot reach down to row k, so that the leg length, l, must be less than $k-1$, so that the equality above cannot hold. Conversely, if $\lambda_k > j-1$ then the hook at $(1,j)$ definitely does reach down, possibly past, row k, so that the leg length l must be at least $k-1$; hence again the equality above cannot hold, and so there can be no repetitions in the sequence above.

We now return to the proof of the lemma: As there are exactly $h_{i,1}$ entries in the hook at box x, and they are all at most $h_{i,1}$, their product must be $(h_{i,1})!$. Therefore

$$(h_{i,1})! \prod_{i<j}(h_{i,1} - h_{j,1}) = \prod_{j=1}^{\lambda_i} h_{i,j}.$$

Finally, the product of all hook lengths is the double product over rows and columns, and gives the required formula.

□

It is this version of the hook length formula that we wish to prove, by induction via removable nodes, so we aim to show that

$$f^\lambda = n!\frac{\prod\limits_{1\leqslant i<j\leqslant r} h_{i,1}-h_{j,1}}{\prod\limits_{i=1}^{r}(h_{i,1})!}.$$

By induction, we have that

$$f^\lambda = \sum_{x\in\mathrm{Rem}(\lambda)}(n-1)!\frac{\prod\limits_{1\leqslant i<j\leqslant r^{(x)}} h^{(x)}_{i,1}-h^{(x)}_{j,1}}{\prod\limits_{i=1}^{r^{(x)}}(h^{(x)}_{i,1})!},$$

where $r^{(x)}$ is the number of parts in $\lambda\setminus\{x\}$ and $h^{(x)}_{i,j}$ are the hook lengths in $\lambda\setminus\{x\}$. If x does not lie in row r, or $\lambda_r>1$, then $r^{(x)}=r$ and $h^{(x)}_{i,j}=h_{i,j}$ except $h^{(x)}_{x,1}=h_{x,1}-1$. If x lies in row r and $\lambda_r=1$ then $r^{(x)}=r-1$ and $h^{(x)}_{i,1}=h_{i,1}-1$ for all i.

Notice that there is at most one removable node per row, the last row always has a removable node, and a row i has a removable node if and only if $h_{i,1}-h_{i+1,1}>1$. In other words, we can replace the formula above with

$$f^\lambda = \sum_{x=1}^{r}(n-1)!\frac{\prod\limits_{1\leqslant i<j\leqslant r^{(x)}} h^{(x)}_{i,1}-h^{(x)}_{j,1}}{\prod\limits_{i=1}^{r^{(x)}}(h^{(x)}_{i,1})!},$$

where the definition of $r^{(x)}$ and $h^{(x)}_{i,j}$ is as above; the extra terms included are all zero, since if there is no removable node in row i then $h^{(x)}_{i,1}-h^{(x)}_{i+1,1}=0$. Our goal is to rearrange this to get something involving only r and $h_{i,1}$: fixing $1\leqslant x\leqslant r$ (assume $\lambda_r>1$), we see that

$$\prod_{i=1}^{r^{(x)}}(h^{(x)}_{i,1})! = \prod_{i=1}^{r}(h_{i,1})!\Big/ h_{x,1},$$

and

$$\prod_{1\leqslant i<j\leqslant r^{(x)}} h^{(x)}_{i,1}-h^{(x)}_{j,1} = \left(\prod_{1\leqslant i<j\leqslant r} h_{i,1}-h_{j,1}\right)\prod_{i>x}\frac{h_{x,1}-1-h_{i,1}}{h_{x,1}-h_{i,1}}\prod_{x>i}\frac{h_{i,1}-h_{x,1}+1}{h_{i,1}-h_{x,1}}.$$

If $\lambda_r = 1$ then both right-hand sides of the two displayed equations differ from the left by a factor of $\prod_{i=1}^{r-1} h_{i,1}$, so their quotient is still correct. Substituting back in, we now get

$$f^\lambda = (n-1)! \frac{\displaystyle\prod_{1 \leqslant i < j \leqslant r} h_{i,1} - h_{j,1}}{\displaystyle\prod_{i=1}^{r} (h_{i,1})!} \cdot \sum_{x=1}^{r} h_{x,1} \cdot \prod_{i \neq x} \left(1 + \frac{1}{h_{i,1} - h_{x,1}}\right).$$

Comparing this formula with the one we want, we need to prove that

$$\sum_{x=1}^{r} h_{x,1} \cdot \prod_{i \neq x} \left(1 + \frac{1}{h_{i,1} - h_{x,1}}\right) = n.$$

In fact, this statement holds for any r distinct complex numbers whose sum is $n + r(r-1)/2$. That is to say, we have the following proposition.

Proposition 1.12 Let $z_1, \ldots, z_r$ be r points in $\mathbb{C}$. We have that

$$\sum_{i=1}^{r} z_i - \frac{r(r-1)}{2} = \sum_{i=1}^{r} z_i \cdot \prod_{j \neq i} \left(1 + \frac{1}{z_j - z_i}\right).$$

Proof: Define a function $f : \mathbb{C} \to \mathbb{C}$ by

$$f(z) = z \prod_{i=1}^{r} \left(1 + \frac{1}{z_i - z}\right),$$

and note that f is holomorphic on $\mathbb{C}$ but not the z_i, with a simple pole at each of the z_i. Thus we can apply Cauchy's residue theorem to get that, for C a circle of large radius at the origin,

$$\int_C f(z) \mathrm{d}z = 2\pi \mathrm{i} \sum_{i=1}^{r} \operatorname{Res}(f, z_i).$$

The residue of the function f at z_i is given by $-z_i \prod_{j \neq i} (1 + 1/(z_j - z_i))$. We now use the residue at infinity; since we have that f is holomorphic on an annulus of infinite outer radius,

$$\frac{-1}{2\pi \mathrm{i}} \int_C f(z) \mathrm{d}z = \operatorname{Res}(f, \infty) = \operatorname{Res}\left(\frac{-1}{z^2} f\left(\frac{1}{z}\right), 0\right).$$

Now we just have to understand $-1/z^2 \cdot f(1/z)$, which is

$$\frac{-1}{z^3} \prod_{i=1}^{r} (1 - z - z_i z^2 + \cdots),$$

and the z^{-1} term of this is $\sum_{i=1}^{r} z_i$ (squares times constant) minus $r(r-1)/2$ (linear times linear). Substituting this in gives the result. $\square$

This now proves the second version of the hook length formula, and so the original version is proved.

Theorem 1.13 (Frame, Robinson, Thrall) If λ is a partition of n, then

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

It turns out that the number of standard tableaux of a given shape is the dimension of a particular representation of the symmetric group S_n. In the next chapter, when we construct all of these, we will associate a $\mathbb{Z}S_n$-module S^λ to any partition λ, and we will see that it has dimension f^λ. These modules, called *Specht modules*, will be the complete set of irreducible representations over the complex field. In particular, this means that

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

However, in order to prove the facts we just stated about Specht modules we will need to know this fact! The next section, on the Robinson–Schensted correspondence, will in particular show this fact (Corollary 1.15).

1.5 The Robinson–Schensted correspondence

The Robinson–Schensted algorithm is another fundamental procedure in the combinatorics of symmetric groups. It assigns to each permutation of n points two standard tableaux of the same shape. We will prove that there is a bijection between permutations of S_n and all pairs of standard tableaux of the same shape, and so the formula

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

in the previous section immediately drops out.

Let $\sigma \in S_n$ be a permutation, sending i to i_σ. We will construct a sequence of 'partial' tableaux (of the same shape)

$$(P_0, Q_0) = (\emptyset, \emptyset),\ (P_1, Q_1), \dots, (P_n, Q_n);$$

at each stage i the number i_σ is inserted into the tableau P_{i-1} at some place, and i is added to Q_{i-1} (in an addable node) to maintain the same shape as P_i. It is clear from what we have just said that Q_i is always standard, and P_i contains numbers 1 to n – but not all of them – and they increase along rows and columns. The procedure, called *row insertion*, is as follows, to insert a number x into a partial tableau P:

(i) Consider the first row of the tableau P: if x is larger than any element of this row, then add x to the end of the row and stop.

(ii) If x is smaller than some element of this row, let y denote the smallest element in the row that is larger than x, replace y by x. (The element y is said to be *bumped* to the next row.)

(iii) Perform the same procedure on the second row of P, with the new value y. Continue until you add a value to the end of a row and stop. (This might be a row that does not exist in P.)

We construct (P_i, Q_i) from (P_{i-1}, Q_{i-1}) by row inserting i_σ into P_{i-1} and adding i to an addable node of Q_{i-1} to make it the same shape as P_i. For example, if $\sigma = (1,6,3)(2,4)$, we get the P_i to be

$$P_0 = \emptyset \quad P_1 = \begin{array}{c} 6 \end{array} \quad P_2 = \begin{array}{c} 4 \\ 6 \end{array} \quad P_3 = \begin{array}{c} 1 \\ 4 \\ 6 \end{array} \quad P_4 = \begin{array}{cc} 1 & 2 \\ 4 & \\ 6 & \end{array} \quad P_5 = \begin{array}{ccc} 1 & 2 & 5 \\ 4 & & \\ 6 & & \end{array} \quad P_6 = \begin{array}{ccc} 1 & 2 & 3 \\ 4 & 5 & \\ 6 & & \end{array}$$

The Q_i are simply

$$Q_0 = \emptyset \quad Q_1 = \begin{array}{c} 1 \end{array} \quad Q_2 = \begin{array}{c} 1 \\ 2 \end{array} \quad Q_3 = \begin{array}{c} 1 \\ 2 \\ 3 \end{array} \quad Q_4 = \begin{array}{cc} 1 & 4 \\ 2 & \\ 3 & \end{array} \quad Q_5 = \begin{array}{ccc} 1 & 4 & 5 \\ 2 & & \\ 3 & & \end{array} \quad Q_6 = \begin{array}{ccc} 1 & 4 & 5 \\ 2 & 6 & \\ 3 & & \end{array}$$

The process of row insertion definitely yields a partial tableau, and so P_n is a standard tableau. We therefore have constructed an algorithm to produce, given any $\sigma \in S_n$, a pair of standard tableau $P_n = P(\sigma)$ and $Q_n = Q(\sigma)$ of the same shape. We now want to reverse the algorithm, taking any pair (P, Q) of standard tableaux of the same shape and returning an element σ of S_n such that $P = P(\sigma)$ and $Q = Q(\sigma)$.

Theorem 1.14 The Robinson–Schensted correspondence is a bijection between permutations and pairs of standard tableaux of the same shape.

Proof: Write (P, Q) for a pair of standard tableaux of the same shape. Let $(P_n, Q_n) = (P, Q)$; we will proceed by induction to construct a sequence (P_i, Q_i) and a permutation σ such that the (P_i, Q_i) are the steps in the Robinson–Schensted algorithm. We prove that we can 'reverse' the algorithm at one particular stage.

The procedure is as follows:

(i) Find the box of Q_n filled by n: this is the removable box of P_n that was added from 'P_{n-1}', filled with x. Delete this box.

(ii) Move to the row above the row that had x in it. To find the box that used to contain x but was bumped by another number, find the largest number y on the row that is smaller than x, put x in the box instead of y and carry y up to the next row.

(iii) Continue this until we run out of rows. The number carried forward this time is σ_n, and the partial tableau that remains is P_{n-1}.

It is clear from our construction that row inserting σ_n into P_{n-1} yields P_n, and so by induction we may construct a sequence (P_i) and a permutation σ such that $P(\sigma) = P$ and $Q(\sigma) = Q$. This completes the proof of the theorem. $\square$

This yields the corollary mentioned twice before, which is explicitly given for future reference.

Corollary 1.15 If n is a positive integer, then

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

Proof: There is a bijection between the permutations of S_n (which total $n!$) and pairs of standard λ-tableaux (which number $(f^\lambda)^2$) as λ ranges over all partitions of n. This yields the above equation. $\square$

1.6 Tabloids and polytabloids

A tableau is a Young diagram with the numbers 1 to n filling them. We can let S_n act on the set of Young tableaux of a fixed shape by permutation of the entries, so that

$$\begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline 4 \\ \cline{1-1} 5 \\ \cline{1-1} \end{array} \cdot (1,5,4)(2,3) = \begin{array}{|c|c|c|}\hline 5 & 3 & 2 \\ \hline 1 \\ \cline{1-1} 4 \\ \cline{1-1} \end{array}.$$

A *tabloid* is where we have a tableau but we ignore the ordering along the rows, so that

$$\begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline 4 & 6 \\ \cline{1-2} 5 \\ \cline{1-1} \end{array} \quad \text{and} \quad \begin{array}{|c|c|c|}\hline 3 & 1 & 2 \\ \hline 6 & 4 \\ \cline{1-2} 5 \\ \cline{1-1} \end{array}$$

are 'the same'. To demonstrate this equality, we remove the vertical bars, as so:

$$\begin{array}{ccc}\hline 1 & 2 & 3 \\ \hline 4 & 6 \\ \cline{1-2} 5 \\ \cline{1-1} \end{array}.$$

If t is a tableau, denote by $\{t\}$ the associated tabloid.

While the action of S_n on tableau was regular (i.e., transitive and the point stabilizer is trivial), the action of S_n on tabloids is transitive but generally far from regular; in the example above the stabilizer is clearly $\mathrm{Sym}(1,2,3) \times \mathrm{Sym}(4,6) \leqslant S_6$. Subgroups of symmetric groups of this type, direct products of symmetric groups acting naturally, are called *Young subgroups*. The isomorphism type is determined completely by the shape of the tabloid, leading to a 'standard' Young subgroup

$$S_\lambda = S_{\lambda_1} \times S_{\lambda_2} \times \cdots \times S_{\lambda_r},$$

associated to a partition $\lambda = (\lambda_1, \ldots, \lambda_r)$, with each S_{λ_i} acting on the next λ_i points.

Given a tableau, the stabilizer of the corresponding tabloid is called the *row stabilizer*, for obvious reasons. Exactly analogously, there is a *column stabilizer*. Denote the row and column stabilizers of a tableau t by R_t and C_t respectively. Combining the two stabilizers results in the concept of a polytabloid, a formal linear combination of tabloids.

Let t be a tableau. The *polytabloid* of t is the formal linear combination

$$\sum_{\sigma \in C_t} \mathrm{sgn}(\sigma)\{t \cdot \sigma\},$$

where $\mathrm{sgn}(\sigma)$ is the sign of the permutation σ and as we said before $\{\cdot\}$ denotes the operation of taking a tabloid of a given tableau. For example, the polytabloid of the tableau

$$\begin{array}{|c|c|c|c|} \hline 1 & 2 & 3 & 4 \\ \hline 5 & 6 \\ \cline{1-2} \end{array}$$

is

$$\begin{array}{cccc} \hline 1 & 2 & 3 & 4 \\ \hline 5 & 6 \\ \cline{1-2} \end{array} - \begin{array}{cccc} \hline 1 & 3 & 4 & 6 \\ \hline 2 & 5 \\ \cline{1-2} \end{array} + \begin{array}{cccc} \hline 3 & 4 & 5 & 6 \\ \hline 1 & 2 \\ \cline{1-2} \end{array} - \begin{array}{cccc} \hline 2 & 3 & 4 & 5 \\ \hline 1 & 6 \\ \cline{1-2} \end{array}.$$

If t is a tableau, denote the associated polytabloid by $\bar{\boldsymbol{t}}$. The element $\sum_{\sigma \in C_t} \mathrm{sgn}(\sigma)\sigma$ will be denoted by κ_t, so that $\bar{\boldsymbol{t}} = \{t\} \cdot \kappa_t$.

Since the symmetric group acts on tabloids it also acts on formal linear combinations of them, so on polytabloids. It also can act on polytabloids by acting on the underlying tableau. We have actions of the symmetric group on many different objects now, and a quick lemma will tell us that these actions are all 'nice'.

Lemma 1.16 If $\sigma \in S_n$, $\lambda \vdash n$ and t is a tableau of shape λ, then we have that $(R_t)^\sigma = R_{t \cdot \sigma}$, $(C_t)^\sigma = C_{t \cdot \sigma}$, and $\kappa_t \cdot \sigma = \kappa_{t \cdot \sigma}$.

The first two parts follow immediately from the fact that R_t is the stabilizer of t as a tabloid and C_t is the stabilizer of t as a 'column tabloid', in which case this is standard theory of transitive permutation groups. That $\kappa_t \cdot \sigma = \kappa_{t\cdot\sigma}$ follows since C_t is transported this way and the sign of a permutation is preserved by conjugation.

We also need to record the action of S_n on the tabloids and polytabloids.

Lemma 1.17 Let $\sigma \in S_n$, t be a tableau of shape λ for some $\lambda \vdash n$. We have that $\{t\} \cdot \sigma = \{t \cdot \sigma\}$ and $\bar{\boldsymbol{t}} \cdot \sigma = \overline{\boldsymbol{t} \cdot \sigma}$.

Proof: To see that $\{t\} \cdot \sigma = \{t \cdot \sigma\}$, notice that the action of taking a tabloid is simply to delete the vertical lines, which obviously commutes with the action of S_n. Finally, since $\bar{\boldsymbol{t}} = \{t\} \cdot \kappa_t$, we have that

$$\bar{\boldsymbol{t}} \cdot \sigma = (\{t\} \cdot \kappa_t) \cdot \sigma = (\{t\} \cdot \sigma) \cdot \kappa_t^\sigma = \{t \cdot \sigma\} \cdot \kappa_{t\cdot\sigma} = \overline{\boldsymbol{t} \cdot \sigma}.$$

$\square$

This proves that S_n acts transitively on the set of tabloids and polytabloids of a given shape.

1.7 Two orderings on partitions

There are two natural orderings on partitions that we will discuss here. The first is standard lexicographic ordering.

Definition 1.18 Let $\lambda = (\lambda_1, \ldots, \lambda_r)$ and $\mu = (\mu_1, \ldots, \mu_s)$ be two partitions of n. The *lexicographic ordering* is defined as follows: we have that $\lambda > \mu$ if, for the first i that $\lambda_i \neq \mu_i$, we have $\lambda_i > \mu_i$.

In this total ordering the partition (n) is the largest element and (1^n) is the smallest element. For the partitions of 4 we have

$$(4) > (3,1) > (2,2) > (2,1,1) > (1^4)$$

and for 5,

$$(5) > (4,1) > (3,2) > (3,1,1) > (2,2,1) > (2,1^3) > (1^5),$$

The other, more subtle, ordering, is a partial ordering called the dominance ordering.

Definition 1.19 Let $\lambda = (\lambda_1, \ldots, \lambda_r)$ and $\mu = (\mu_1, \ldots, \mu_s)$ be two partitions of n. The *dominance ordering* is defined as follows: we have that $\lambda \trianglerighteq \mu$ if for all $1 \leqslant n \leqslant r$, we have

$$\sum_{i=1}^{n} \lambda_i \geqslant \sum_{i=1}^{n} \mu_i.$$

The dominance ordering is certainly a partial ordering rather than a total ordering, but again it has a unique maximal element of (n), and a unique minimal element of (1^n). For the partitions of 4 and 5, the dominance ordering is the same as the total ordering, but for $n = 6$,

$$(6) > (5,1) > (4,2) > (4,1,1) > (3,3) > (3,2,1) > (3,1^3) > (2,2,2) > (2,2,1,1) > (2,1^4) > (1^6),$$

whereas $(4,1,1)$ and $(3,3)$ are incomparable under $\trianglerighteq$. Chains that match up with $>$ are

$$(6) \trianglerighteq (5,1) \trianglerighteq (4,2) \trianglerighteq (4,1,1), \quad (3,3) \trianglerighteq (3,2,1) \trianglerighteq (3,1^3), \quad (2,2,2) \trianglerighteq (2,2,1,1) \trianglerighteq (2,1^4) \trianglerighteq (1^6),$$

and we can connect the three chains via $(4,1,1) \trianglerighteq (3,2,1)$ and $(4,2) \trianglerighteq (3,3)$ for the former two, and $(3,1^3) \trianglerighteq (2,2,1,1)$ and $(3,2,1) \trianglerighteq (2,2,2)$ for the latter two.

The following obvious lemma, suggested in the examples above, relates the partial dominance ordering to the total lexicographic ordering.

Lemma 1.20 If $\lambda \trianglelefteq \mu$ then $\lambda \leqslant \mu$.

We need one lemma on the dominance ordering in the next chapter.

Lemma 1.21 (Dominance lemma) Let t be a tableau of shape λ and let s be a tableau of shape μ. If, for every index i, the elements of row i of s lie in different columns of t, then $\mu \trianglelefteq \lambda$.

Proof: We will relabel the entries $1, \dots, n$ of t, preserving the columns of t: we can do this so that the entries in the first i rows of s can all be placed in the first i rows of t. (To see that this is possible, since for each row the entries in s appear in different columns of t, there are at most i entries from the first i rows of s in each column of t.) Doing this, we see that the number of entries in the first i rows of t is at least that of the first i rows of s, so that $\mu \trianglelefteq \lambda$, as claimed. $\square$

1.8 An ordering on tabloids

Here we will construct an analogue of the dominance ordering on partitions, but this time for tabloids of a fixed shape λ. For our purposes – proving that polytabloids associated to standard λ-tableaux are linearly independent – there are several different partial orderings that will work. The key property is the following: if t is a standard tableau, then for any $\sigma \in C_t$, $\{t\}$ dominates $\{t \cdot \sigma\}$. Since polytabloids are formal linear combinations of tabloids it makes sense to take linear combinations of polytabloids. To see that the property we highlighted just now is the appropriate one, we prove the result.

Proposition 1.22 Suppose that $\sum_i a_i \overline{\boldsymbol{t}_i} = 0$ is a linear combination of standard λ-polytabloids $\overline{\boldsymbol{t}_i}$ with $a_i \in R$ and R a commutative ring. If $\trianglelefteqslant$ is a partial ordering on tabloids such that if t is a standard tableau then for any $\sigma \in C_t$, $\{t\} \trianglerighteqslant \{t \cdot \sigma\}$, then all a_i are zero.

Proof: Order the t_i so that whenever $\{t_i\} \trianglelefteqslant \{t_j\}$ we have $i \geqslant j$. Notice that $\{t_i\}$ appears with coefficient a_i in $a_i\overline{\boldsymbol{t}_i}$. We claim that $\{t_1\}$ does not appear in any other $\overline{\boldsymbol{t}_i}$, so that since the sum is zero, a_1 must be zero. By induction then, all a_i are zero. If $\{t_1\}$ appears in $\overline{\boldsymbol{t}_i}$ then by the fundamental property of $\trianglelefteqslant$, $t_1 \trianglerighteqslant t_i$. This contradicts the ordering on the t_i, and so the result holds. $\square$

In order to prove that the $\overline{\boldsymbol{t}_i}$ are linearly independent then, it suffices to prove the existence of such an ordering $\trianglelefteqslant$. There are several choices for such an ordering: one is given in Exercise 2.1, and we give a simpler one here. If $\{t\}$ is a tabloid, let t be the tableau with increasing rows whose tabloid in $\{t\}$. The *column word* of such a tabloid is the sequence of integers between 1 to n, starting in the left-most column of t and reading from bottom to top. We write $\{t\} \trianglerighteqslant \{s\}$ if the largest number, i.e., n where t has size n, appears earlier in t than in s, and if n appears in the same place then consider $n-1$, and so on. If the column words of s and t are equal, then write $\{t\} \trianglerighteqslant \{s\}$ if $\{t\} = \{s\}$. For example, if $\{t\}$ is the tabloid

$$\begin{array}{ccccc} \hline 7 & 5 & 3 & 8 & 1 \\ \hline 9 & 4 & 2 & & \\ \cline{1-3} 6 & 10 & & & \\ \cline{1-2} \end{array},$$

then the column word is $6, 2, 1, 10, 4, 3, 9, 5, 7, 8$, and the maximal element has position 4.

Notice that there are many tabloids with the same column word, but if one fixes the column word and the shape then one fixes the tabloid. Also, if $\{t\}$ is a tabloid and $\sigma \in S_n$, then either $\{t \cdot \sigma\} = \{t\}$ or the column words of $\{t\}$ and $\{t \cdot \sigma\}$ are different, in which case $\{t\}$ and $\{t \cdot \sigma\}$ are comparable under $\trianglelefteqslant$.

If t is a standard tableau then the column word of $\{t\}$ is the column word of t, and along each column the entries increase down the column; if $\sigma \in C_t$ and i is the largest integer moved by σ, then i appears further up the column of $\{t \cdot \sigma\}$. Hence all integers greater than i appear in the same place in the column words of $\{t\}$ and $\{t \cdot \sigma\}$, and i appears earlier in $\{t\}$ than in $\{t \cdot \sigma\}$. Thus if t is standard then $\{t\} \trianglerighteqslant \{t \cdot \sigma\}$, exactly the property required.

We need to use this result in a different context, so we rewrite it in terms of general vector spaces.

Proposition 1.23 Let $e_1, \dots, e_n$ be a basis for a finite-dimensional vector space V, and let $\trianglerighteqslant$ be a partial order on the elements of V. Suppose that $v_1, \dots, v_m$ are elements of V such that the v_i each contain a unique maximal element m_i under $\trianglerighteq$, and the m_i are all distinct, then the v_i are all linearly independent.

Chapter 2

Constructing the representations of the symmetric group

2.1 M^λ and S^λ

The first construction is of the M^λ, the ambient module within which we will construct the S^λ.

Definition 2.1 Let λ be a partition of n. The *Young module*, denoted M^λ, is the permutation module of S_n with basis the set of tabloids of shape λ.

This is of course just the permutation module on the cosets of the Young subgroup S_λ. This definition is independent of the field, or even ring, over which we want to do our representation theory. Traditionally a lot of these definitions are done over $\mathbb{Z}$, since this encapsulates both the modular theory (the field $\mathbb{F}_p$) and the ordinary theory (the field $\mathbb{C}$).

The construction of S^λ is almost as easy.

Definition 2.2 Let λ be a partition of n. The *Specht module*, denoted S^λ, is the submodule of M^λ generated by the polytabloids of shape λ.

A *cyclic module* is a module M such that there exists $m \in M$ with $M = \langle m \rangle$ (as a module). We see that, since S_n acts transitively on tabloids and polytabloids, both Young and Specht modules are cyclic, generated by any tabloid and polytabloid respectively.

Unlike the case of M^λ, the generating set for S^λ is not a basis. We already know the dimension of S^λ – it should be f^λ, and so one good guess for a basis of S^λ is the set of polytabloids $\bar{\boldsymbol{t}}$ where t is a *standard* tableau of shape λ. Indeed, this is the case, but this will take some time to prove. First, we need to prove that Specht modules are irreducible, at least over $\mathbb{C}$, which we do now.

Let $(\cdot,\cdot)$ denote the unique bilinear form such that $(\{t\},\{s\}) = \delta_{\{t\},\{s\}}$. (Since the tabloids form a basis for M^λ, such a bilinear form exists.) We first prove that if t is a λ-tableau and s is a μ-tableau, then $\{s\}\cdot\kappa_t = 0$ unless λ dominates μ, and if $\lambda = \mu$ then $\{s\}\cdot\kappa_t$ is a (possibly zero) scalar multiple of $\{t\}$.

Suppose that $\{s\}\cdot\kappa_t \neq 0$. If a and b are two labels in the same row of s that also lie in the same column of t, then $(a,b)\in C_t$ and we therefore see that $\kappa_t = (1-(a,b))x$, where x consists of all the even permutations in C_t by Exercise 2.2. As $\{s\}(a,b) = \{s\}$ by hypothesis,

$$\{s\}\cdot\kappa_t = \{s\}\cdot(1-(a,b))x = 0.$$

Thus no two entries in the same row of s lie in the same column of t, and firstly λ dominates μ by the dominance lemma, and moreover the argument in the proof of the dominance lemma states that there exists $\sigma\in C_t$ such that $\{t\}\cdot\sigma = \{s\}$. We therefore get

$$\{s\}\cdot\kappa_t = (\{t\}\cdot\sigma)\kappa_t = \operatorname{sgn}(\sigma)(\{t\}\cdot\kappa_t) = \pm\bar{\boldsymbol{t}}.$$

The equality $\sigma\kappa_t = \kappa_t\sigma = \operatorname{sgn}(\sigma)\kappa_t$ when $\sigma\in C_t$ is Exercise 2.2.

Theorem 2.3 (James's submodule theorem) Let V be a submodule of M^λ. Either V contains S^λ or $V\subset (S^\lambda)^\perp$.

Proof: By the above paragraph, if t and s are two λ-tableaux then $\{s\}\cdot\kappa_t$ is a scalar multiple of $\bar{\boldsymbol{t}}$. We therefore see that if m is any element of M^λ and t is a tableau of shape λ, then $m\kappa_t$ is a scalar multiple of $\bar{\boldsymbol{t}}$, simply because m is expressible as a linear combination of λ-tableaux, and we apply the displayed equation above.

Suppose that there exists $v\in V$ and t a tableau of shape λ such that $v\cdot\kappa_t\neq 0$. As V is a submodule, $v\cdot\kappa_t\in V$, but $v\cdot\kappa_t$ is a scalar multiple of $\bar{\boldsymbol{t}}$, so $\bar{\boldsymbol{t}}\in V$. This proves that $S^\lambda \leqslant V$ since $S^\lambda = \langle\bar{\boldsymbol{t}}\rangle$.

Conversely, we could have that, for all $v\in V$ and t a tableau of shape λ, $v\cdot\kappa_t = 0$. Exercise 2.2 now tells us that, given any λ-tableau t and $v\in V$, we have that

$$(v,\bar{\boldsymbol{t}}) = (v, t\cdot\kappa_t) = (v\cdot\kappa_t, t) = 0.$$

Consequently, $V\leqslant (S^\lambda)^\perp$ (since the $\bar{\boldsymbol{t}}$ span S^λ as a vector space), as claimed. $\square$

All of this was done over any field, in fact over any ring R where we think of RG-modules as free over R (technically these are *RG-lattices*, but since we will be solely interested in fields, this distinction need not be made). We see the following corollary.

Corollary 2.4 If λ is a partition of n then $S^\lambda/(S^\lambda\cap(S^\lambda)^\perp)$ is either 0 or irreducible.

Over a field of characteristic 0, since our bilinear form is clearly non-degenerate, we see that $S^\lambda \cap (S^\lambda)^\perp = 0$, so that the Specht modules are irreducible. We now prove that if $\lambda \neq \mu$ then S^λ and S^μ are non-isomorphic. This will unfortunately require the field to be the complex numbers for now, although we will later repair this deficiency and prove that the S^λ form a complete set of irreducible kS_n-modules for any field k of characteristic 0 or $p > n$.

Theorem 2.5 If $\mathrm{Hom}(S^\lambda, M^\mu)$ is non-zero then λ dominates μ, and $\mathrm{Hom}(S^\lambda, M^\lambda)$ consists of multiplication by a scalar. Consequently, the modules S^λ for $\lambda \vdash n$ form a complete set of irreducible kS_n-modules for any field k of characteristic 0.

Proof: Let $\phi \in \mathrm{Hom}(S^\lambda, M^\mu)$ be non-zero, so that there exists some λ-tableau t such that $\bar{\boldsymbol{t}}\phi \neq 0$. As $\mathrm{char}\, k = 0$, $M^\lambda = S^\lambda \oplus (S^\lambda)^\perp$, so let $\bar\phi \in \mathrm{Hom}(M^\lambda, M^\mu)$ be given by extending ϕ via projection along S^λ first. Hence

$$0 \neq (\{t\} \cdot \kappa_t)\phi = (\{t\}\phi)\kappa_t = \left(\sum_{\{s\}} c_s\{s\}\right)\kappa_t = \sum_s c_s(\{s\} \cdot \kappa_t).$$

In the paragraph before James's submodule theorem we proved that if $\{s\} \cdot \kappa_t \neq 0$ then λ dominates μ, where t has shape λ and s has shape μ. The displayed equation above proves that in our case λ dominates μ as one of the $\{s\} \cdot \kappa_t$ must be non-zero. Furthermore, if $\lambda = \mu$ then $\{s\} \cdot \kappa_t$ is either 0 or a scalar multiple of $\{t\}$, so that $\bar{\boldsymbol{t}}\phi = c\bar{\boldsymbol{t}}$ for some (non-zero) scalar c. As S^λ is generated by $\bar{\boldsymbol{t}}$, we have that ϕ is scalar multiplication by c, as needed.

To see that the Specht modules form a complete set of irreducible kS_n-modules, firstly there are the correct number of them, namely equal to the number of partitions of n, so it suffices to show that they are pairwise non-isomorphic. Suppose that $S^\lambda = S^\mu$; then in particular $\mathrm{Hom}(S^\lambda, M^\mu)$ and $\mathrm{Hom}(S^\mu, M^\lambda)$ are non-zero, so that λ dominates μ and vice versa, yielding $\lambda = \mu$, and this completes the proof. $\square$

This also shows that the permutation module M^μ decomposes as a sum of $m_{\lambda,\mu}S^\lambda$ for various $\lambda \trianglerighteq \mu$, with $m_{\lambda,\lambda} = 1$. The $m_{\lambda,\mu}$ are called *Kostka numbers*, and have a purely combinatorial description in terms of so-called 'semistandard' tableaux, which we will see later in this chapter.

We end with a result that we promised earlier.

Theorem 2.6 Over any commutative ring, a basis for S^λ consists of the standard λ-polytabloids, and $\dim S^\lambda = f^\lambda$.

Proof: The standard λ-polytabloids are linearly independent by Proposition 1.22, so $\dim S^\lambda \geqslant f^\lambda$. However, by the Robinson–Schensted correspondence $n! = \sum_\lambda (f^\lambda)^2$, and since

$$n! = |S_n| = \sum_{\lambda \vdash n} \dim(S^\lambda)^2 \geqslant \sum_{\lambda \vdash n} (f^\lambda)^2 = n!,$$

we see that the polytabloids $\bar{\boldsymbol{t}}$ for t a standard λ-tableau are a basis for S^λ, yielding $\dim S^\lambda = f^\lambda$. $\square$

2.2 The branching rule

What happens when you restrict the $\mathbb{C}S_n$-module S^λ from S_n to S_{n-1}? In the inductive proof of the hook length formula, we obtained the formula

$$f^\lambda = \sum_{x \in \mathrm{Rem}(\lambda)} f^{\lambda \setminus \{x\}},$$

and since $f^\lambda = \dim S^\lambda$, the obvious guess would be

$$S^\lambda = \bigoplus_{x \in \mathrm{Rem}(\lambda)} S^{\lambda \setminus \{x\}}.$$

The obvious guess turns out to be correct. Since the dimensions of both sides agree, and Specht modules are all irreducible, it suffices to show that for every $x \in \mathrm{Rem}(\lambda)$, the module $S^{\lambda \setminus \{x\}}$ appears as a composition factor in the restriction of S^λ, for then by Maschke's theorem we are done.

In order to prove this we will construct a tower of submodules whose successive quotients will, as S_{n-1}-modules, consist of $S^{\lambda \setminus \{x\}}$, for x starting at the uppermost removable node and proceeding down λ.

Since there is at most one removable node in a given row, we can label the rows with removable nodes $r_1, \dots, r_m$, with $r_i < r_{i+1}$. If t is a standard tableau, then the label n must lie in one of the rows r_i. Let V_i be the subspace of S^λ spanned by standard polytabloids $\bar{\boldsymbol{t}}$ where the label n appears in the first r_i rows of t. We claim that V_i is an S_{n-1}-submodule of S^λ and that V_i/V_{i-1} is isomorphic as an S_{n-1}-module to $S^{\lambda^{(i)}}$, where $\lambda^{(i)}$ is obtained from λ by removing the last box on row r_i. This then proves, at least if $\mathrm{char}\, p = 0$ or $\mathrm{char}\, p > n$, the direct sum decomposition of the restriction of S^λ as described above.

Since S_{n-1} fixes the box containing n, it permutes the polytabloids with n in row at most r_i; hence the subspaces V_i are also S_{n-1}-submodules of S^λ. Construct a linear map $\phi_i : M^\lambda \to M^{\lambda^{(i)}}$ by specifying $\{t\}\phi = 0$ unless the label n appears in the r_ith row, in which case $\{t\}\phi$ is $\{t\}$ with the label n removed. Since S_{n-1} preserves the location of the label n, its action commutes with the removal of the label n, so ϕ_i is a $\mathbb{C}S_{n-1}$-module homomorphism. If t is a standard tableau with n in the r_ith row then, if s lies in the C_t-orbit of t, the label n appears in row r_i if and only if any element of C_t mapping t to s lies in $C_t \cap S_{n-1}$, and is also true if and only if $s\phi_i$ is non-zero. Hence we get

$$\overline{\boldsymbol{t}}\phi_i = (\{t\}\kappa_t)\phi_i = \sum_{\sigma \in C_t \cap S_{n-1}} \operatorname{sgn}(\sigma)\{t'\} \cdot \sigma = \overline{\boldsymbol{t}'},$$

where t' denotes the tableau obtained from $\overline{\boldsymbol{t}}$ by removing the label n. If t is a standard tableau with n in the r_ith row then t' is a standard tableau of shape $\lambda^{(i)}$, so the image of S^λ under ϕ_i, which contains $\overline{\boldsymbol{t}'}$, must contain all of $S^{\lambda^{(i)}}$ since $S^{\lambda^{(i)}}$ is generated as an S_{n-1}-module by $\overline{\boldsymbol{t}'}$. Thus $S^\lambda\phi_i$ contains $S^{\lambda^{(i)}}$, and we have proved that S^λ has each $S^{\lambda^{(i)}}$ as a composition factor. This completes the proof of the branching rule, by our observation at the start.

(In fact, it is possible to prove without much more effort that $V_i/V_{i-1} \cong S^{\lambda^{(i)}}$, since we notice that $S^{\lambda^{(i)}}$ dominates $S^{\lambda^{(j)}}$ whenever $j < i$. We do not need this, however.)

We now use Frobenius reciprocity to get the full branching rule.

Theorem 2.7 (Branching rule) Let λ be a partition of n. We have that, as $\mathbb{C}S_n$-modules,

(i) $S^\lambda \downarrow_{S_{n-1}} = \bigoplus_{x \in \mathrm{Rem}(\lambda)} S^{\lambda \setminus \{x\}}$ and

(ii) $S^\lambda \uparrow^{S_{n+1}} = \bigoplus_{x \in \mathrm{Add}(\lambda)} S^{\lambda \cup \{x\}}$.

Proof: The first part has been proved above. To see the second one, we use Frobenius reciprocity; this states that if H is a subgroup of G, M is a $\mathbb{C}G$-module and N is a $\mathbb{C}H$-module, then $\mathrm{Hom}(M \downarrow_H, N) = \mathrm{Hom}(M, N \uparrow^G)$. Specializing to the case $G = S_{n+1}$, $H = S_n$, $M = S^\mu$ for $\mu \vdash (n+1)$, $N = S^\lambda$ for $\lambda \vdash n$, we see that $\mathrm{Hom}(S^\mu, S^\lambda \uparrow^{S_{n+1}})$ is equal to $\mathrm{Hom}(S_\mu \downarrow_{S_n}, S^\lambda)$. This latter space is 1-dimensional if $\lambda = \mu \setminus \{x\}$ for some removable node x, or equivalently $\mu = \lambda \cup \{x\}$ for some addable node x, so $S^\lambda \uparrow^{S_{n+1}}$ contains a single copy of $S^{\lambda \cup \{x\}}$ for each $x \in \mathrm{Add}(\lambda)$, and nothing else. Maschke's theorem now tells us that this is a direct sum, yielding (ii). $\square$

2.3 The characteristic p case

In Corollary 2.4 we said that $S^\lambda/(S^\lambda \cap (S^\lambda)^\perp)$ is either 0 or irreducible. If it is non-zero, call this module D^λ. It turns out that, as in the characteristic 0 case, the D^λ yield all irreducible representations of S_n, and furthermore that no two of them are isomorphic. However, in the modular case (i.e., characteristic p where $p \leqslant n$) we will have fewer irreducible representations of D_n than partitions, so for some partitions λ we must have that $S^\lambda \leqslant (S^\lambda)^\perp$, so that there is no D^λ. We saw this already in Exercise 2.4 with the partition (1^n) and $p \leqslant n$; we wish to generalize this and come up with a necessary and sufficient condition for $S^\lambda \leqslant (S^\lambda)^\perp$.

We start with proving one direction of this, that if there are p parts of a partition λ that have the same length then $S^\lambda \leqslant (S^\lambda)^\perp$.

Proposition 2.8 Let t and t' be λ-tableaux. Write a_j for the number of parts of λ of size j. We have that $(\overline{\boldsymbol{t}}, \overline{\boldsymbol{t}'})$ is a multiple of $\prod_j a_j!$. In particular, if k is a field of characteristic p and $a_j \geqslant p$ for some j, then $(\ ,\ )$ is the zero form on S^λ (viewed as a kS_n-module), so that $S^\lambda \leqslant (S^\lambda)^\perp$.

Proof: Let X denote the set of all λ-tabloids that appear in both $\overline{\boldsymbol{t}}$ and $\overline{\boldsymbol{t}'}$. Firstly note that if $\{u\}$ is a λ-tabloid in X, and $\{v\}$ can be obtained from $\{u\}$ by reordering the rows, then $\{v\} \in X$: there is definitely a permutation in $\sigma \in C_t$ that reorders the rows to send $\{u\}$ to $\{v\}$ (and therefore also $\tau \in C_{t'}$), and moreover since the sign of this permutation is simply determined by the action on the rows, $\mathrm{sgn}(\sigma) = \mathrm{sgn}(\tau)$. This implies that the product of the signs with which $\{u\}$ appears in both $\overline{\boldsymbol{t}}$ and $\overline{\boldsymbol{t}}'$ is the same as that for $\{v\}$. Thus X is a union of equivalence classes X_i under this action of row reordering, and furthermore the products of the coefficients in front of each element of a fixed X_i from $\overline{\boldsymbol{t}}$ and $\overline{\boldsymbol{t}'}$ are the same, so the contribution of X_i to $(\overline{\boldsymbol{t}}, \overline{\boldsymbol{t}'})$ is $\pm|X_i|$. As $|X_i| = \prod a_j!$, we get the result.

Finally, if one of the a_j is at least p then p divides the result of evaluating the bilinear form on every pair of generating elements for S^λ, so over k of characteristic p this is the zero form, and $S^\lambda \leqslant (S^\lambda)^\perp$. $\square$

We say that a partition λ of n is *t-regular* if there are at most $t-1$ parts of any given size in λ. We see here that the if λ is not p-regular then there is no simple module D^λ in characteristic p. Of course, we are now tasked with proving the converse, namely that $D^\lambda \neq 0$ if λ *is* p-regular.

We have proved that the form evaluated on two polytabloids is a multiple of $\prod a_j!$. We now find two λ-tableaux, t and t^*, for which the form evaluates to a small multiple of this product. Since this is non-zero in a field of characteristic p whenever λ is p-regular, this proves that $D^\lambda > 0$, and completes the proof. (We then of course need to show that no two D^λ are isomorphic, and that these are all irreducibles over a field of characteristic p.)

Proposition 2.9 Let λ be a partition of n, and suppose that it has exactly a_j parts of size j. If t is a λ-tableau, and t^* is the λ-tableau obtained by reversing the entries in the rows of

t, then

$$(\overline{\boldsymbol{t}}, \overline{\boldsymbol{t}^*}) = \prod_j (a_j!)^j.$$

In particular, if k is a field of characteristic p and $a_j < p$ for all j, then $(\ ,\ )$ is not the zero form on S^λ (viewed as a kS_n-module), so that $S^\lambda \nleq (S^\lambda)^\perp$.

Proof: Suppose that $\{t\sigma\} = \{t^*\sigma^*\}$ is a tabloid appearing in both $\overline{\boldsymbol{t}}$ and $\overline{\boldsymbol{t}^*}$ (so that $\sigma \in C_t$ and $\sigma^* \in C_{t^*}$). We claim that $\sigma = \sigma^*$. Without loss of generality, suppose that the $(1, 1)$ entry of t is 1: then the $(1, \lambda_1)$ entry of t^* is also 1, and σ^*, stabilizing column λ_1, can only send 1 into a row that is also of size λ_1.

Hence if there are b_1 parts of λ that have size λ_1, then the last entries in these b_1 rows of t^* (and also of t) must be permuted. Since $\{t\sigma\} = \{t^*\sigma^*\}$, $t\sigma$ and $t^*\sigma^*$ have the same rows, so that σ and σ^* must agree on the set of entries in the first and last column of the first b_1 rows of t.

Now consider the next b_2 rows that have the same size, directly below the first b_1 rows. The last entries of these rows of t and t^* can only be sent to the first $b_1 + b_2$ rows by elements of C_t and C_{t^*}, and the row that σ and σ^* send a given entry must be the same since $t\sigma$ and $t^*\sigma^*$ have the same tabloid. However, the first entries of the first b_1 rows are fixed, so in fact the last entries of these rows of t and t^* must be permuted amongst themselves; hence σ and σ^* agree on these elements by the same argument as for the first b_1 rows.

Repeating this argument proves that σ and σ^* agree on the entries in the first column of both t and t^*. We now remove the first and last columns of t and t^* and apply induction on the number of columns of a tableau to see that $\sigma = \sigma^*$. What we also see that is that σ sends entries in a given row i only to rows of the same length as i.

Hence this element σ actually lies in $C_t \cap C_{t^*}$. This intersection consists of exactly those permutations that stabilize the union of the rows of a given length, the property of σ in the previous paragraph. We therefore see that the set of tabloids appearing in both $\overline{\boldsymbol{t}}$ and $\overline{\boldsymbol{t}^*}$ is

$$\{\{t\sigma\} \mid \sigma \in C_t \cap C_{t^*}\}.$$

Thus

$$(\overline{\boldsymbol{t}}, \overline{\boldsymbol{t}^*}) = \sum_{\sigma \in C_t \cap C_{t^*}} (\operatorname{sgn}(\sigma) \cdot \{t\sigma\}, \operatorname{sgn}(\sigma) \cdot \{t^*\sigma\}) = |C_t \cap C_{t^*}|.$$

Finally, we see that $C_t \cap C_{t^*}$ stabilizes the blocks consisting of rows of equal size, and thus it acts as the symmetric group of size a_j on each of the j columns in each block, so each block constributes $(a_j!)^j$. Hence

$$|C_t \cap C_{t^*}| = \prod_j (a_j!)^j,$$

as required. $\square$

Having done this, we now isolate a bit about which D^λ can appear in a given S^μ in characteristic p.

Proposition 2.10 Let k be a field of characteristic p. Suppose that $\lambda, \mu \vdash n$, and that λ is p-regular. If V is any submodule of M^μ, and there exists a non-zero homomorphism $\phi : S^\lambda \to M^\mu/V$ then λ dominates μ. If $\lambda = \mu$ and t is a λ-tableau, then $\bar{\boldsymbol{t}}\phi \in \langle \bar{\boldsymbol{t}} + V \rangle$. Consequently, if D^λ and D^μ are isomorphic then $\lambda = \mu$.

Proof: Suppose that λ has a_j parts of size j, so that $\sum ja_j = n$. Let t be any λ-tableau, and write t^* for the tableau obtained from t by reversing the entries in each row, as in Proposition 2.9. Just before Theorem 2.3 we proved that $\{s\}\kappa_t$ is a scalar multiple of $\bar{\boldsymbol{t}}$ for all λ-tableaux s, so $\overline{\boldsymbol{t}^*}\kappa_t = \gamma\bar{\boldsymbol{t}}$ for some γ. Utilizing Exercise 2.2 and Proposition 2.9, and the fact that $(\bar{\boldsymbol{t}}, \{t\}) = 1$ we compute

$$\gamma = \gamma(\bar{\boldsymbol{t}}, \{t\}) = (\overline{\boldsymbol{t}^*}\kappa_t, \{t\}) = (\bar{\boldsymbol{t}}^*, \{t\}\kappa_t) = (\bar{\boldsymbol{t}}^*, \bar{\boldsymbol{t}}) = \prod_{j=1}^{n} (a_j!)^j.$$

In particular, γ is not the zero element of the field k of characteristic p. The Specht module S^λ is a cyclic module, generated by $\bar{\boldsymbol{t}}$, and ϕ is a non-zero homomorphism, so $\bar{\boldsymbol{t}}\phi$ must be non-zero, so

$$\overline{\boldsymbol{t}^*}\phi\kappa_t = \gamma\bar{\boldsymbol{t}}\phi \neq 0;$$

we see that $\overline{\boldsymbol{t}^*}\phi = v$ has a non-zero product with κ_t. Since M^μ has basis the μ-tabloids, there must be a μ-tabloid $\{r\}$ such that $\{r\} + V$ appears in v and $v\kappa_t \neq 0$. Again, above Theorem 2.3 we proved that this implies that $\lambda \trianglerighteq \mu$. Furthermore, if $\lambda = \mu$ then $M^\lambda\kappa_t = \langle \bar{\boldsymbol{t}} \rangle$, so we get the result.

Finally, if D^λ and D^μ are isomorphic then $\lambda \trianglerighteq \mu$ and $\mu \trianglerighteq \lambda$, so $\lambda = \mu$, completing the proof. $\square$

Thus any composition factor D^λ of M^μ satisfies $\lambda \trianglerighteq \mu$. Write $d_{\lambda,\mu}$ for the number of copies of D^μ in S^λ. The previous proposition has the following corollary.

Corollary 2.11 Ordering the rows of the matrix $D = (d_{\lambda,\mu})$ by lexicographic order on the p-regular partitions first, then the partitions that are not p-regular, and the lexicographic order on the columns, we see that D is lower trinagular.

The matrix D is the *decomposition matrix*, and basically measures how simple modules over the complex numbers decompose into irreducibles when viewed over a field of characteristic p. Drawing a graph with vertices the S^λ and connecting two vertices S^λ and S^μ if there exists ν such that $d_{\lambda,\nu}$ and $d_{\mu,\nu}$ are both non-zero (so that there is a column that has non-zero entries for both S^λ and S^μ at the same time in the decomposition matrix), the connected components of this matrix are the *blocks* of S_n. (This graph and definition can be made for all groups.) The connection between cores and the representation theory can be made now.

Theorem 2.12 (Nakayama conjecture) Let λ and μ be partitions of n. Two Specht modules S^λ and S^μ lie in the same p-block if and only if λ and μ have the same p-core.

The proof of this theorem is beyond the scope of this course, but it places a further restriction on the decomposition numbers $d_{\lambda,\mu}$.

The final thing we need to do is prove that any simple kS_n-module is one of the D^λ. Brauer proved that the number of simple modules over an algebraically closed field of characteristic p is the number of conjugacy classes of *p-regular* elements, i.e., elements or order prime to p. For symmetric groups, this is conjugacy classes of elements without a pm-cycle for some m, or equivalently the number of partitions whose parts have size prime to p. Hence we need the following result.

Theorem 2.13 (Glaisher's theorem) Let n and d be positive integers. The number of partitions of n for which there are at most $d-1$ parts of any given size is equal to the number of partitions whose parts have size not divisible by d.

Proof: Write $\mathcal{A}_n$ for the set of all partitions n with at most $d-1$ parts of any given size, and write $\mathcal{B}_n$ for the set of partitions of n whose parts have size not divisible by d. Let λ be a partition of n, with exactly a_j parts of size j. We will define two mutually inverse functions between $\mathcal{A}_n$ and $\mathcal{B}_n$.

Let $\phi : \mathcal{A}_n \to \mathcal{B}_n$ be defined by the following: if there is a part λ_i of $\lambda \in \mathcal{A}_n$ of size divisible by d, replace λ_i by d parts of size λ_i/d. This (eventually) results in a partition of n without any parts divisible by d, so ϕ has the right codomain.

Let $\psi : \mathcal{B}_n \to \mathcal{A}_n$ be defined by the following: if there are d parts of size λ_i, replace them by one part of size $d\lambda_i$. This (eventually) results in a partition of n with at most $d-1$ parts of any given size, so ψ has the right codomain.

It is easy to see that ϕ and ψ invert each other, so these two sets have the same cardinality. □

Corollary 2.14 The modules D^λ for λ a p-regular partition of n form a complete set of the simple kS_n-modules, for any field k of characteristic p.

2.4 Young's rule

The permutation module M^μ can be written, over the complex field certainly, as a sum of Specht modules. We now that $\mathrm{Hom}(S^\lambda, M^\mu)$ is non-zero, i.e., S^λ appears as a summand of M^μ, only if λ dominates μ, so we would like to know the multiplicities $m_{\lambda,\mu}$, the Kostka numbers as they were called earlier in this chapter.

Young's rule determines these Kostka numbers. To compute them we need the concept of a semistandard tableau, which is like a standard tableau although the numbers may repeat. If μ is a composition (with zeroes allowed) of n then a *tableau of type* μ is a Young diagram with entries in its n boxes, with μ_i copies of the number i. For example, a Young tableau of type $(0, 2, 0, 1, 1)$ is the tableau

$$\begin{array}{|c|c|}\hline 2 & 4 \\ \hline 5 & 2 \\ \hline \end{array}\ .$$

If μ is the composition (1^n) then a Young tableau of type μ is simply a Young tableau. We will write $\underline{t}$ for a tableau whose type is not necessarily (1^n), reserving symbols like t for genuine tableaux. A Young tableau of type μ is *semistandard* if the entries weakly increase along the rows and strictly increase along the columns. We generally take μ itself to be a partition, in which case we get the following lemma.

Lemma 2.15 Let λ and μ be partitions of n. If there exists a semistandard partition of shape λ and type μ then λ dominates μ, and there is exactly one semistandard tableau of shape λ and type λ.

Proof: If $\underline{t}$ is a semistandard λ-tableau then any number i appears in at most the ith row of $\underline{t}$. Hence, since there are $\mu_1 + \cdots + \mu_i$ numbers that are at most i in $\underline{t}$,

$$\lambda_1 + \lambda_2 + \cdots + \lambda_i \geqslant \mu_1 + \mu_2 + \cdots + \mu_i;$$

thus $\lambda \trianglerighteq \mu$. If $\mu = \lambda$, we see that there is a unique way of placing $\lambda_1 + \cdots + \lambda_i$ copies of i in the first i rows for all i, namely row i consists solely of is. Hence there is a single semistandard tableau of shape λ and type λ. $\square$

That looks like the conditions that $m_{\lambda,\mu} > 0$ implies λ dominates μ and $m_{\lambda,\lambda} = 1$. In fact, this is not a coincidence.

Theorem 2.16 (Young's rule) The Kostka number $m_{\lambda,\mu}$ is the number $K_{\lambda,\mu}$ of semistandard tableaux of shape λ and type μ.

To prove this theorem we need an alternative definition of the permutation module M^μ. Let $T_{\lambda,\mu}$ denote the set of all tableaux of shape λ and type μ. The first thing to notice is that there are the same number of these as there are tabloids of shape μ. To see this, clearly the shape of λ is irrelevant, as we can place the rows of λ one after another in a single row, so that λ becomes the partition (n). In this case, there are $n!$ different ways of placing the entries, but all 1s are alike, as are 2s and so on, giving $n!/\prod_j a_j!$ (where μ has a_j parts of size j), the same as the number of μ-tabloids.

Hence the vector space spanned by formal linear combinations of tableaux of shape λ and type μ is isomorphic as a vector space to M^μ. To make them isomorphic as kS_n-modules we must transport the S_n action to the vector space $kT_{\lambda,\mu}$. To do this, fix a λ-tableau t, and for any tableau $\underline{s}$ of shape λ of any type, write $\underline{s}_i$ for the entry of d in the position labelled by i in the fixed tableau t. We could for simplicity fix t to be the standard λ-tableau that uses the integers 1 to n in order from left to right then top to bottom, so for $\lambda = (3,3,2)$ we get

$$\begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline 4 & 5 & 6 \\ \hline 7 & 8 & \multicolumn{1}{c}{} \\ \cline{1-2}\end{array}\,.$$

In which case, if $\underline{s}$ is the λ-tableau

$$\begin{array}{|c|c|c|}\hline 7 & 4 & 2 \\ \hline 1 & 1 & 4 \\ \hline 2 & 3 & \multicolumn{1}{c}{} \\ \cline{1-2}\end{array}$$

then $\underline{s}_4 = \underline{s}_5 = 1$ and $\underline{s}_1 = 7$.

Given any tabloid $\{s\}$ of shape μ, produce a tableau $\underline{s}$ of shape λ and type μ by setting $\underline{s}_i$ to be the row in which i appears in $\{s\}$, so that if s is the tableau

$$\begin{array}{|c|c|c|c|}\hline 8 & 1 & 5 & 2 \\ \hline 3 & 7 & \multicolumn{2}{c}{} \\ \cline{1-2} 6 & \multicolumn{3}{c}{} \\ \cline{1-1} 4 & \multicolumn{3}{c}{} \\ \cline{1-1}\end{array}$$

of shape $(4,2,1,1)$ then $\underline{s}$ is the tableau

$$\begin{array}{|c|c|c|}\hline 1 & 1 & 2 \\ \hline 4 & 1 & 3 \\ \hline 2 & 1 & \multicolumn{1}{c}{} \\ \cline{1-2}\end{array}\,.$$

of shape λ and type μ. We get that this map $\phi : \{s\} \mapsto \underline{s}$ is a bijection between the μ-tabloids and $T_{\lambda,\mu}$. We define the action of S_n on $T_{\lambda,\mu}$ to be 'the same' as on the set of μ-tabloids, i.e., we want that ϕ is an isomorphism of kS_n-modules. For this to hold, we need for $\{s\}$ a μ-tabloid to have

$$(\{s\}\cdot\sigma)\phi = (\{s\}\phi)\cdot\sigma = \underline{s}\cdot\sigma.$$

To do this, we actually have to define $(\underline{s} \cdot \sigma)_i$ to be $\underline{s}_{i\sigma^{-1}}$. To see that this is the right definition, we see that

$$\begin{aligned}(\underline{s} \cdot \sigma)_i &= \text{ row number of } i\sigma^{-1} \text{ in } \{s\} \\ &= \text{ row number of } i \text{ in } \{s \cdot \sigma\} \\ &= \underline{s \cdot \sigma_i}\end{aligned}$$

This proves that ϕ is a module homomorphism, essentially because we defined it to be so.

For a tableau $\underline{s}$ of shape λ and type μ, define $\phi_{\underline{s}}$ to be the map sending the fixed tableau t to the sum of all members of $T_{\lambda,\mu}$ that are row equivalent to $\underline{s}$, then extend this to a map $\phi_{\underline{s}} : M^\lambda \to kT_{\lambda,\mu}$ by the fact that $\{t\}$ generated the kS_n-module M^λ. (Notice that this map is well defined since the stabilizer of $\{t\}$ under the action of S_n is the row stabilizer R_t, and by definition of the action of S_n on $T_{\lambda,\mu}$, this permutes the rows of $\underline{s}$, hence preserves the sum of all members of $T_{\lambda,\mu}$ that are row equivalent to $\underline{s}$.

We may also restrict this action to S^λ, to get homomorphisms $\phi_{\underline{s}} : S^\lambda \to M^\mu$. Some of these will be zero, and for $k = \mathbb{C}$ at least, the rest will be isomorphisms. (If k has characteristic p then there might be maps that are neither 0 nor isomorphisms.) We claim that, whenever $\underline{s}$ is semistandard, $\phi_{\underline{s}}$ is non-zero, and also that these maps between vector spaces are linearly independent. To prove this latter statement, we need the following obvious lemma.

Lemma 2.17 Let V and W be vector spaces, and let $\phi_1, \ldots, \phi_n$ be a set of linear maps $\phi_i : V \to W$. If there exists $v \in V$ such that the $v\phi_i$ are linearly independent in W, then the ϕ_i are linearly independent.

Let us now prove that they are linearly independent. As with the Specht module, we prove linear independence using a partial ordering on the semistandard tableaux, and a result that looks exactly like Proposition 1.22. Write $\{\underline{s}\}$ for the row equivalence class in $T_{\lambda,\mu}$ containing $\underline{s}$.

We have that

$$(\bar{\boldsymbol{t}})\phi_{\underline{s}} = (\{t\}\kappa_t)\phi_{\underline{s}} = \left(\sum_{\underline{u} \in \{\underline{s}\}} \underline{u}\right)\kappa_t.$$

This consists of tableau in $T_{\lambda,\mu}$ obtained from $\underline{s}$ by first permuting the rows and then permuting the columns. The appropriate column word we need now, since the rows can also be permuted, is to read from the last column to the first, reading from the bottom to the top, so the column word of

1	4	5	2
4	4	3	
1	6		

becomes $2,3,5,6,4,4,1,4,1$. The appropriate partial order we need is the following: for $\underline{s}, \underline{u} \in T_{\lambda,\mu}$, and j is the largest integer such that μ_j is non-zero, we write the column words for $\underline{s}$ and $\underline{u}$, and declare $\underline{s} \trianglerighteq \underline{u}$ if for all i, the ith occurrence of an entry μ_j is earlier in $\underline{s}$ then in $\underline{u}$ (or of course, $\underline{s} = \underline{u}$); if the ns are in the same place, we move to the $(n-1)$s and so on, as with the partial order on tabloids.

The next easy lemma gives the property of $\trianglerighteq$ that we need.

Lemma 2.18 If $\underline{s} \in T_{\lambda,\mu}$ is semistandard and $\underline{u}$ appears in the expression $\bar{\boldsymbol{t}}\phi_{\boldsymbol{s}}$, then $\underline{s} \trianglerighteq \underline{u}$.

Proof: Let j be the largest integer such that $\mu_j > 0$. The entries of $\underline{s}$ that are μ_j must be at the end of their columns and of their rows, by the semistandard nature of $\underline{s}$. For each particular entry x that is labelled μ_j, applying a row equivalence must move it to an earlier column (or swap it with another μ_j, which can be ignored), and then any column equivalence results in it still appearing later than it does in $\underline{s}$. Hence all occurrences of μ_j in the column word of $\underline{s}$ are earlier than those in $\underline{u}$. $\square$

We can now apply Proposition 1.23. In this case, the v_i are the $\bar{\boldsymbol{t}}\phi_{\underline{s}}$ for $\underline{s}$ semistandard and the m_i are the $\underline{s}$ themselves. Hence the $\bar{\boldsymbol{t}}\phi_{\underline{s}}$ are all linearly independent and therefore the $\phi_{\underline{s}}$ are.

We claim that $K_{\lambda,\mu} = K_{\lambda,\mu}$ from Exercise 2.5. As with the proof that the standard λ-tableaux are a basis for S^λ, we need only count dimensions. Exercise 2.5(iv) states that

$$\sum_{\lambda \vdash n} K_{\lambda,\mu} \dim S^\lambda = \dim M^\mu,$$

so that since we know that there are *at least* $K_{\lambda,\mu}$ copies of S^λ lying in M^μ then we are done. This completes the proof of Young's rule.

Chapter 3

The Murnaghan–Nakayama rule

The Murnaghan–Nakayama rule is a combinatorial algorithm to compute the character values of any character χ^λ of a Specht module S^λ on any conjugacy class of cycle type μ. In order to prove this rule, we will need to firstly deal with the Littlewood–Richardson rule, which computes the multiplicity of χ^μ in the induced character $(\chi^\lambda \otimes \chi^\mu)\uparrow^{S_n}$, where λ is a partition of m, μ is a partition of $n-m$, and this tensor product is thought of as a character for $S_m \times S_{n-m}$, viewed as a Young subgroup of S_n.

We start with skew tableaux, which are needed for the Murnaghan–Nakayam rule, and then the Littlewood–Richardson rule, which also needs these skew tableaux.

3.1 Skew tableaux

Skew partitions and skew tableaux are objects obtained by removing boxes from a Young diagram to form a snake-like object.

Definition 3.1 Let λ and μ be partitions, with $\mu_i \leqslant \lambda_i$ for all i. Write $\lambda \setminus \mu$ for the object obtained by removing all boxes from λ that are also present in μ. This is called *skew partition* of shape $\lambda \setminus \mu$. If the boxes of $\lambda \setminus \mu$ are filled with numbers, we call this a *skew tableau* of shape $\lambda \setminus \mu$ and type ν in the same way as for tableaux. If the entries of a skew tableau are weakly increasing along rows and strictly increasing down columns then it is *semistandard*.

Definition 3.2 A *lattice word* is a sequence $a_1, \dots, a_m$ from the integers $1, \dots, n$, such that, for every initial subword of $a_1, \dots, a_m$, for every i the number of a_j that are equal to i is at least as many as there are equal to $i+1$.

In other words, a lattice word is one where, if one draws a histogram with $1, \dots, n$ on the x-axis and the number of a_j equal to a given i on the y-axis, then as more of the a_j are added to the total, the columns of the histogram never increase reading from left to right.

A lattice word is sometimes called a *ballot word* for the following voting interpretation: if the candidates are labelled $1, \ldots, n$, then as the ballots are counted, candidate i is always beating or tied with candidate $i+1$ for all i.

An example of a lattice word is 1123213, and an example of a word that is not a lattice word is 11123132.

Definition 3.3 The *row-reversed word* of a skew tableau is the word obtained by reading each row from right to left, from the top row to the bottom row. A *Littlewood–Richardson skew tableau* is a skew semistandard tableau for which the row-reversed word is a lattice word.

While there are always Littlewood–Richardson tableaux of a given shape $\nu \setminus \lambda$, if one also fixes the type μ then sometimes there are some, and sometimes there are none. It is easy to see that in a Littlewood–Richardson tableau the first row consists solely of 1s, so that is a first condition. Write $c^{\nu}_{\lambda,\mu}$ for the number of Littlewood–Richardson tableaux of shape $\nu \setminus \lambda$ and type μ. We always assume that $|\lambda| + |\mu| = |\nu|$ in what follows.

The following lemma looks trivial, but in fact describes the branching rule of the previous chapter, although it will take some time to prove this.

Lemma 3.4 Suppose that $|\nu| = |\lambda| + 1$, and that $\mu = (1)$. We have that $c^{\nu}_{\lambda,\mu} = 1$ if $\nu = \lambda \cup \{x\}$ for some $x \in \mathrm{Add}(\lambda)$, and $c^{\nu}_{\lambda,\mu} = 0$ otherwise.

In a similar vein, this next lemma is Young's rule for $K_{\lambda,\mu}$, where μ is a two-part partition.

Lemma 3.5 Suppose that $\lambda = (m)$ and $\mu = (n-m)$. We have that $c^{\nu}_{\lambda,\mu}$ is equal to 1 if $\nu = (n-a, a)$ for $a \leqslant m$, and 0 otherwise. This coincides with the number of semistandard tableaux of shape ν and type $(n-m, m)$.

Again the proof is omitted, but what these two lemmas are trying to say is that the coefficients $c^{\nu}_{\lambda,\mu}$ are describing the multiplicity of the character χ^{ν} in some $\chi^{\lambda} \otimes \chi^{\mu}$ induced up to S_n. This will be proved in the next section.

3.2 The Littlewood–Richardson rule

The branching rule computed the consistuents of $\chi^{\lambda} \uparrow^{S_n}$, where $\lambda \vdash (n-1)$. Viewing S_{n-1} as a Young subgroup $S_{n-1} \times S_1$, corresponding to the partition $(n-1, 1)$, the branching rule really computes the multiplicity of χ^{μ} (for $\mu \vdash n$) in the induced character $(\chi^{\lambda} \otimes \chi^{(1)}) \uparrow^{S_n}$, where $\chi^{\lambda} \otimes \chi^{(1)}$ is viewed as a character on the direct product $S_{n-1} \times S_1$.

At the same time, Young's rule computes the multiplicity of S^λ in the permutation module M^μ. One may think of M^μ as the induced module from the trivial module for $S_{\mu_1} \times \cdots \times S_{\mu_s}$, and so Young's rule actually computes the decomposition into irreducible characters of the permutation character

$$(\chi^{(\mu_1)} \otimes \chi^{(\mu_2)} \otimes \cdots \otimes \chi^{(\mu_s)}) \uparrow^{S_n} .$$

The Littlewood–Richardson rule simultaneously generalizes both of these, by computing the characters of any irreducible character, not just the trivial, induced from any Young subgroup, not just S_{n-1}.

Theorem 3.6 (Littlewood–Richardson rule) We have that

$$(\chi^\lambda \otimes \chi^\mu) \uparrow^{S_n} = \sum_{\nu \vdash n} c^\nu_{\lambda,\mu} \chi^\nu.$$

In order to prove this theorem we need to extend our situation to a more general one. Let Λ denote the set of all partitions of all integers (including 0), and let $\mathbb{Q}\Lambda$ denote the vector space of all (finite) linear combinations of elements of Λ. (This is really nothing more than a notational device, somewhere to make our statements.) Similar to lattice words, if $a_1, \ldots, a_m$ is any word, with entries from $1, \ldots, n$, we say that a_j is *good* if either $a_j = 1$, or in the initial sequence $a_1, \ldots, a_j$, the number of good terms equal to a_j is at most the number of good terms equal to $a_j - 1$. (Notice that this definition is inductive.) The *quality* of a word $a_1, \ldots, a_m$, is $\alpha_1, \ldots, \alpha_n$, where α_i is the number of good terms in the word equal to i. A word is a lattice word if and only if all terms are good.

We introduce operators $[\mu^*, \mu]$ on $\mathbb{Q}\Lambda$ for μ a composition and μ^* a partition with $\mu^*_i \leqslant \mu_i$ in a similar way to lattice words: define

$$\lambda^{[\mu^*,\mu]} = \sum a_\nu \nu,$$

where $a_\nu = 0$ unless $\lambda_i \leqslant \nu_i$ for all i (i.e., the skew diagram $\nu \setminus \lambda$ exists) and in this case it is equal to the number of skew semistandard tableaux of shape $\nu \setminus \lambda$ and type μ whose row-reversed words has quality at least μ^* (i.e., the quality $\beta_1, \ldots, \beta_r$ satisfies $\beta_i \geqslant \mu^*_i$). Since a word is a lattice word if and only if all terms are good, a skew semistandard tableau of type μ is a Littlewood–Richardson tableau if and only if its quality is μ. Hence the operator $[\mu, \mu]$ applied to λ yields

$$\lambda^{[\mu,\mu]} = \sum c^\nu_{\lambda,\mu} \nu.$$

Statements we can prove about the operator $[\mu^*, \mu]$ will specialize to statements about $c^\nu_{\lambda,\mu}$ by setting $\mu^* = \mu$, so we have indeed extended our situation.

The two operators of interest are $[0, \nu]$ and $[\nu, \nu]$, and we will want to be able to move between them easily. Most of our preparation will be in understanding these operators and their effect.

Lemma 3.7 If $\mu = (\mu_1, \ldots, \mu_s)$ is a composition of n, then $(0)^{[0,\mu]} = \sum_{\nu \vdash n} K_{\nu,\mu}\nu$, i.e., the constituents of the permutation module M^μ. If μ is a partition, then $0^{[\mu,\mu]} = \mu$.

Proof: Every sequence is of quality $0, \ldots, 0$, so in this case we are counting the (skew) semistandard tableaux of shape $\nu \setminus \emptyset$ and type μ, i.e., $K_{\nu,\mu}$, as claimed.

For the second part of the lemma, we need to understand the number of semistandard tableaux of shape $\nu \vdash n$ and weight μ whose row-reversed word is a lattice word. Firstly, placing entry i into row i yields a semistandard tableau of shape μ and weight μ with the required property, and furthermore it is the only semistandard tableau of shape μ and type μ, so μ appears in $0^{[\mu,\mu]}$ with coefficient 1. We also know that if $K_{\lambda,\mu} > 0$ then $\lambda \trianglerighteq \mu$, so let s be a semistandard tableau of shape λ and type μ. If $\lambda_1 > \mu_1$ then there exist integers greater than 1 in the first row of s, hence its row-reversed word is not a lattice word. Hence $\lambda_1 = \mu_1$, and by the same argument applied repeatedly we get $\lambda_i = \mu_i$ for all i, as the first i for which $\lambda_i \neq \mu_i$ we must have $\lambda_i > \mu_i$ since $\lambda \trianglerighteq \mu$. This completes the proof. $\square$

We now examine the operator $[0, \mu]$ a bit more closely, where $\mu = (\mu_1, \ldots, \mu_r)$. This requires us to add all possible skew semistandard tableaux of type μ onto a given partition λ, but with no condition on the row-reversed word. In other words, this is the same as sequentially adding the μ_1 boxes with entry 1, then the μ_2 boxes with entry μ_2, the only condition being the semistandardness of the resulting skew tableau, which is the same as simply requiring semistandardness of each addition. Hence we get the equivalence of operators

$$[0, \mu] \equiv [\mu_1, \mu_1][\mu_2, \mu_2] \ldots [\mu_r, \mu_r].$$

The previous lemma, together with this, proves the following.

Corollary 3.8 Let $\mu = (\mu_1, \ldots, \mu_r)$ be a composition of n. We have that

$$0^{[\mu_1,\mu_1]\ldots[\mu_r,\mu_r]} = \sum_{\nu \vdash n} K_{\nu,\mu}\nu.$$

The last bit of information we need is to do with adding and raising operators. This is only needed for one simple fact, and the proof is complex, so we will omit it, but struggle with it in the exercise sheet instead. (Yippee.)

Let μ^* and μ be as above, and suppose that $\mu^* \neq \mu$. Let $i > 1$ be such that $\mu^*_i < \mu_i$ but $\mu^*_{i-1} = \mu_{i-1}$.

We define the adding operator to be as follows: $\mu^* A_i, \mu$ is the pair $0, 0$ if $\mu^*_{i-1} = \mu^*_i$, and if not then $\mu^* A_i$ differs from μ^* in that the ith component has 1 added to it. Notice that $\mu^* A_i$ remains a partition if μ^* is.

We define the raising operator as follows: $\mu^*, \mu R_i$ is the pair of compositions obtained from μ^* and μ by replacing μ_i by μ_i^* and replacing μ_{i-1} by $\mu_{i-1} + \mu_i - \mu_i^*$. In other words, this shifts boxes from the ith row to the $(i-1)$th row. The compositions μ and μR_i are both compositions of the same integer, and we still have that $\mu_j^* \leqslant (\mu R_i)_j$ for all j.

These adding and raising actions affect the operators $[\mu^*, \mu]$ in a nice way.

Proposition 3.9 As operators, whenever $\mu_{i-1}^* = \mu_{i-1}$ but $\mu_i^* < \mu_i$ we have

$$[\mu^*, \mu] \equiv [\mu^* A_i, \mu] + [\mu^*, \mu R_i].$$

Write $\lambda \circ \mu$ for the sum $\sum_\nu c_{\lambda,\mu}^\nu \nu$, so that (with the obvious meaning) $(\chi^\lambda \otimes \chi^\mu) \uparrow^{S_n} = \chi^{\lambda \circ \mu}$.

Using Proposition 3.9, we may write, for any partition $\nu \vdash n$, the operator $[0, \nu]$ as a sum of operators $c_\gamma[\gamma, \gamma]$ for $\gamma \vdash n$, where each c_γ is a non-negative integer, $c_\nu = 1$, and $c_\gamma = 0$ unless γ dominates ν. To see this, firstly note that $[0, \nu] = [\nu_1, \nu]$, since any 1 is good, so we can apply A_2 and R_2 to $[\nu_1, \nu]$ and get us started. Firstly, A_i always increases $|\nu^*|$, so induction on that means that it can never decrease. Secondly, νR_i yields a composition higher up the dominance ordering than ν, so induction on the dominance ordering proves that we eventually end with pairs of partitions $[\gamma, \gamma]$, with $\gamma \trianglerighteq \nu$. To see quickly that $c_\nu = 1$, use Lemma 3.7.

Thus the matrix of such coefficients c_γ is unitriangular, hence invertible. This yields expressions

$$[\lambda, \lambda] = \sum_\alpha a_\alpha [0, \alpha], \qquad [\mu, \mu] = \sum_\beta b_\beta [0, \beta].$$

We will use these expressions to derive the result. We calculate as follows:

$$\begin{aligned}
\lambda^{[\mu,\mu]} &= \left(0^{[\lambda,\lambda]}\right)^{[\mu,\mu]} \\
&= \left(0^{\sum a_\alpha [0,\alpha]}\right)^{\sum b_\beta [0,\beta]} \\
&= 0^{\sum_{\alpha,\beta} a_\alpha b_\beta [\alpha_1,\alpha_1] \ldots [\alpha_a,\alpha_a][\beta_1,\beta_1]\ldots[\beta_b,\beta_b]} \\
&= \sum_{\alpha,\beta} a_\alpha b_\beta \sum_\nu K_{\nu,\alpha\cup\beta} \nu.
\end{aligned}$$

Since $(M^\lambda \otimes M^\mu) \uparrow^{S_n} = M^{\lambda \cup \mu}$, this expression, which is a sum of permutation characters of $M^{\alpha\cup\beta}$, can be factorized into:

$$\begin{aligned}
&= \left(\sum_\alpha a_\alpha \sum_\lambda K_{\lambda,\alpha} \lambda\right) \left(\sum_\beta b_\beta \sum_\nu K_{\nu,\beta} \nu\right) \\
&= \left(0^{\sum a_\alpha [0,\alpha]}\right) \circ \left(0^{\sum b_\beta [0,\beta]}\right) \\
&= 0^{[\lambda,\lambda]} \circ 0^{[\mu,\mu]} \\
&= \lambda \circ \mu.
\end{aligned}$$

This proves that $\lambda^{[\mu,\mu]} = \lambda \circ \mu$, and this is the statement of the Littlewood–Richardson rule, completing its proof.

Corollary 3.10 For any partitions λ and μ, $\lambda^{[\mu,\mu]} = \mu^{[\lambda,\lambda]}$, so $c^{\nu}_{\lambda,\mu} = c^{\nu}_{\mu,\lambda}$.

3.3 The Murnaghan–Nakayama rule

In this section we use the Littlewood–Richardson rule to determine character values of χ^λ on a permutation σ of cycle type μ. Suppose that $\mu \neq (n)$; then (choosing a conjugate of σ) we can write $\sigma = \pi\rho$, where $\pi = (1, \dots, m)$ has cycle type (μ_1) and ρ has cycle type $\nu = (\mu_2, \dots, \mu_r)$. Notice that $\pi\rho$ lies in the subgroup $S_m \times S_{n-m}$, so this suggests that we can use the Littlewood–Richardson rule and Frobenius reciprocity. We see that

$$\chi^\lambda(\pi\rho) = \chi^\lambda \downarrow_{S_m \times S_{n-m}} (\pi\rho) = \sum_{\substack{\alpha \vdash m \\ \beta \vdash n-m}} a^{\lambda}_{\alpha,\beta} \chi^\alpha(\pi)\chi^\beta(\rho),$$

where $a^\lambda_{\alpha,\beta}$ are some positive integers. However, Frobenius reciprocity says that there should be a connection between $\chi^\lambda \downarrow_{S_m \times S_{n-m}}$ and $(\chi^\alpha \otimes \chi^\beta) \uparrow^{S_n}$; this is in fact that

$$a^\lambda_{\alpha,\beta} = \langle \chi_\lambda \downarrow_{S_m \times S_{n-m}}, \chi^\alpha \otimes \chi^\beta \rangle = \langle \chi_\lambda, (\chi^\alpha \otimes \chi^\beta) \uparrow^{S_n} \rangle = c^\lambda_{\alpha,\beta},$$

so the $a^\lambda_{\alpha,\beta}$ are really just Littlewood–Richardson coefficients. So we have accomplished something here. We have

$$\chi^\lambda(\sigma) = \sum_{\beta \vdash n-m} \chi^\beta(\rho) \sum_{\alpha \vdash m} c^\lambda_{\alpha,\beta} \chi^\alpha(\pi).$$

We have two things to sort out now: the first is to evaluate $\chi^\alpha(\pi)$, which is this case of σ being an n-cycle again, and computing $c^\lambda_{\alpha,\beta}$. Then we will have a recursive description of $\chi^\lambda(\sigma)$, since β has fewer parts than λ.

We can put off the n-cycle case no longer, and so here it is.

Proposition 3.11 Let $\sigma \in S_n$ be an n-cycle, and let $\lambda \vdash n$. We have that

$$\chi^\lambda(\sigma) = \begin{cases} (-1)^\ell & \lambda = (n-\ell, 1^\ell), \\ 0 & \text{otherwise.} \end{cases}$$

Proof: Let $0 < m < n$ be any integer, and let α be a partition of m and β be a partition of $n-m$. We claim that if either α or β satisfies $\alpha_2 > 1$ or $\beta_2 > 1$, then so does every partition supported in $\alpha^{[\beta,\beta]}$. If $\alpha_2 > 1$ then this clear, and since $\alpha^{[\beta,\beta]} = \beta^{[\alpha,\alpha]}$ the result holds. If $\alpha_2, \beta_2 \leqslant 1$ – such partitions are called *hook partitions* – then we next claim that $\alpha^{[\beta,\beta]}$ is supported on only two hook partitions, namely $(a+b, 1^{n-a-b})$ and $(a+b-1, 1^{(n-a-b+1)})$. To see this, we need to add a Littlewood–Richardson skew semistandard tableau of weight $\beta = (b, 1^{n-m-b})$ to $\alpha = (a, 1^{m-a})$. The only way we can do this is to add either all 1s to the first row of α and then $2, 3, \ldots, n-m-b+1$ to the bottom of the first column, or all but one of the 1s to the first row of α, and then $1, 2, 3, \ldots, n-m-b+1$ to the bottom of the first column. This proves the claim.

Finally, we get to the proof of the proposition. We see that the inner product of $\chi^{\alpha\circ\beta}$ with the alternating sum of all χ^γ for γ a hook partition, i.e.,

$$\left\langle \chi^{\alpha\circ\beta}, \sum_{c=1}^{n} (-1)^{n-c} \chi^{(c,1^{n-c})} \right\rangle = 0,$$

since either $\alpha \circ \beta$ has either no hook partitions at all or two with differing signs. Moreover, letting ϕ denote the alternating sum of characters of hook partitions, we see that since $\langle \chi^{\alpha\circ\beta}, \phi \rangle = 0$, the same holds for $\langle \chi^\alpha \otimes \chi^\beta, \phi \downarrow_{S_m \times S_{n-m}} \rangle$, so that $\phi \downarrow_{S_m \times S_{n-m}}$ has no constituents at all, i.e., is the zero character, and this holds for *any* $0 < m < n$. Therefore $\phi(\tau) = 0$ unless τ is an n-cycle, since else $\tau \in S_m \times S_{n-m}$ for some m.

Now consider the character table of S_n, and let $\boldsymbol{v}$ denote the column vector consisting of $(-1)^{n-c}$ for the character $\chi^{(c,1^{n-c})}$ and 0 for all other characters, so that this is the representation of ϕ in the character table. We have shown that, apart from the class of n-cycles, this column vector is orthogonal to all other columns, so since the character table of any finite group is invertible (Exercise 3.2) this column must be a scalar multiple of the column of n-cycles. However, we know that $\chi^{(n)}(1) = 1$, and this is the entry for the first row of $\boldsymbol{v}$ as well, so that $\boldsymbol{v}$ must be *exactly* the column of character values on n-cycles. This completes the proof. $\square$

We therefore have reduced the expression for $\chi^\lambda(\pi\rho)$ that was given before the proposition to

$$\chi^\lambda(\sigma) = \sum_{\beta \vdash n-m} \chi^\beta(\rho) \sum_{a=0}^{m-1} c^\lambda_{(a,1^{m-a}),\beta} (-1)^{m-a}.$$

The last thing to do is evaluate $c^\lambda_{(a,1^{m-a}),\beta} = c^\lambda_{\beta,(a,1^{m-a})}$, and then do a little bit of algebra to get the final form.

Definition 3.12 Let λ be a partition of n. The *rim* of λ is the collection of boxes $x = (i, j)$ such that there is no box $(i+1, j+1)$ in λ. A *rim hook* is a connected subcollection ξ of the rim such that $\lambda \setminus \xi$ is a partition.

If ξ is a hook then there is a corresponding rim hook, which consists of the part of the rim between the ends of the arm and leg of ξ, and this action is a bijection between rim hooks and hooks. This allows us to talk about rim t-hooks, leg lengths $ll(\xi)$ of rim hooks, and so on. Removing a rim hook is the same as removing the corresponding hook, so we have not done anything really new. For example, the rim of $(5, 5, 3, 3, 2)$ is the following:

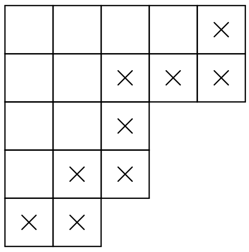

The projection of a hook onto the rim is easy to visualize:

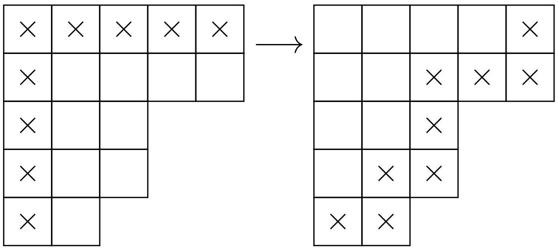

Proposition 3.13 Let β be a partition of $n - m$, λ be a partition of n, and write $\alpha = (a, 1^{m-a})$. If $c^{\lambda}_{\beta,\alpha} > 0$ then $\lambda \setminus \beta$ is a disconnected union of rim hooks. If there are k components spanning c columns, then

$$c^{\lambda}_{\beta,\alpha} = \binom{k-1}{c-a}.$$

Proof: Since $\alpha = (a, 1^{m-a})$, the entries in the skew tableau $\lambda \setminus \beta$ are a copies of 1 and the numbers $2, \ldots, m - a + 1$. Firstly we show that there can be no two-by-two subsquare in the skew tableau $\lambda \setminus \beta$. To see this, column strictness means the bottom-left and bottom-right entries of the subsquare are greater than 1, say i and j respectively, then semistandardness tells us that $i < j$. This however contradicts the lattice word property, as in the row-reversed word, j appears before i (and there is only one copy of each). Hence there are no two-by-two squares, so $\lambda \setminus \beta$ is a subset of the rim of λ. This proves the first part of the proposition. Let k denote the number of conencted components, and let ξ be such a component.

Running down any column of the component ξ, only the top entry can be a 1 by column strictness. Note that each integer greater than 1 appears exactly once, so in the row-reversed word, i must occur before $i+1$, so $i+1$ must occur either below or to the left of i. However, semistandardness requires that $i + 1$ appears to the right of i if they are in the same row, so no two integers greater than 1 appear in the same row. Thus every row consists either entirely of 1s or of 1s followed by an integer greater than 1. This means that the only flexibility about which integers go where in ξ is whether the top-right box consists of a 1 or not. Notice that since the top of every column of ξ is a 1, apart possibly from the far-right column, if ξ spans c' columns then we must use $c'-1$ copies of 1.

We now count how many 1s we have left once we use them across the column tops of the components. If there are c columns spanned by k components, we saw above that we are forced to place $c-k$ copies of 1 there. This leaves $a-(c-k)$ copies of 1 for the k components. However, the top component much have a 1 at the start, by the lattice word condition, so there are $a-c+k-1$ to distribute amonst $k-1$ components. This gives

$$\binom{k-1}{a-c+k-1} = \binom{k-1}{c-a}$$

ways. □

We now need a binomial fact: we have that

$$\sum_{a=0}^{b} (-1)^a \binom{b}{a} = \delta_{b,0}.$$

Using this, we get that

$$\sum_{a=0}^{m-1} c^{\lambda}_{(a,1^{m-1}),\beta}(-1)^{m-a} = \sum_{a=0}^{m-1} (-1)^{m-a} \binom{k-1}{c-a}.$$

Since the binomial expression is 0 outside of the range $0 \leqslant c-a < k$, i.e., $c-k < a \leqslant c$, which values definitely appear in the sum, so this sum becomes

$$= \sum_{a=c-k+1}^{c} (-1)^{m-a} \binom{k-1}{c-a} = \begin{cases} (-1)^{m-c} & k=1, \\ 0 & \text{otherwise.} \end{cases}$$

The statement $k=1$ means that $\lambda \setminus \beta$ is a single skew hook with m squares and c columns, so that $m-c$ is the leg length of $\lambda \setminus \beta$. Putting all this together, we get that

$$\chi^{\lambda}(\pi\rho) = \sum_{|\xi=m|} (-1)^{ll(\xi)} \chi^{\lambda\setminus\xi}(\rho).$$

This gives us the Murnaghan–Nakayama rule.

Theorem 3.14 (Murnaghan–Nakayama rule) Let λ be a partition of n, and let $\mu = (\mu_1, \dots, \mu_r)$ be a composition of n. If σ has cycle type μ and $\tau \in S_{n-\mu_1}$ has cycle type $(\mu_2, \dots, \mu_r)$, then

$$\chi^{\lambda}(\sigma) = \sum_{\xi} (-1)^{ll(\xi)} \chi^{\lambda\setminus\xi}(\tau),$$

where the sum runs over all rim μ_1-hooks of λ.