End of Semester Report
URA Fall 2006

Brendan Fry

The purpose of studying representations of finite groups is to better understand abstract groups. We will be studying the representations of the symmetric group S_n, and how they relate to the partitions of n. More specifically, we look at how we can use p-irreducible partitions to help determine the irreducible representations of S_n.

S_n is a group whose elements are the permutations π of the set $\{1, \ldots, n\}$. These elements can be written as the product of disjoint cycles, and the *cycle type* of π is an expression of the form

$$(1^{m_1}, 2^{m_2}, ..., n^{m_n}),$$

where m_i is the number of cycles of length i in π. We can also give the cycle type in terms of a *partition of n*. This is simply a sequence

$$\lambda = (\lambda_1, ..., \lambda_l),$$

where the λ_i's are weakly decreasing and such that $\sum_{i=1}^{l} \lambda_i = n$.

What is interesting about these partitions and cycle types is their relation to the conjugacy classes of S_n. The set of all elements conjugate to an element π in S_n is called the *conjugacy class of* π. Since conjugacy is an equivalence relation, the distinct conjugacy classes partition S_n. One shows that two elements in S_n are in the same conjugacy class if and only if they have the same cycle type. Thus we get a natural bijective correspondence between the partitions of n and the conjugacy classes of S_n.

Since we want to look at the partitions of n to establish the irreducible representations of S_n, we must first see how to construct the Specht modules which determine these representations. To do so, we need these definitions:

Definition 0.1 Let V be a vector space over a field F and G be a group. Then V is a *G-module* or *module* if there is a group homomorphism

$$\rho : G \to GL(V),$$

where *GL(V)* is the set of all invertible linear transformations of V to itself.

Definition 0.2 Let V be a *G-module.* A *submodule of* V is a subspace W that is closed under the action of G, i.e.,

$$\mathbf{w} \in W \Rightarrow g\mathbf{w} \in W \ \ \forall g \in G.$$

Definition 0.3 A nonzero G-module V is *reducible* if it contains a non-trivial submodule W. V is called *irreducible* if it not reducible. Intuitively, one should think of an irreducible module as one that cannot be broken down into smaller modules.

The following definitions help to create the Specht module, a specific type of module that we must use in studying the representations of S_n.

Notation. If $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_l)$ is a partition of n, then we write $\lambda \vdash n$.

Definition 0.4 Suppose $\lambda \vdash n$. The *Ferrers diagram*, or *shape*, of λ is an array of n boxes having l left-justified rows with row i containing λ_i dots for $1 \leq i \leq l$.

EXAMPLE. If our partition $\lambda = (5, 5, 3, 2, 1)$, then the corresponding Ferrers diagram of λ is

$$\begin{array}{ccccc} \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & & \\ \bullet & \bullet & & & \\ \bullet & & & & \end{array} .$$

Definition 0.5 A *Young tableau of shape* λ is an array t obtained by replacing the dots of the Ferrers diagram of λ with the numbers $1, 2, \ldots, n$ bijectively.

EXAMPLE. If our partition $\lambda = (2, 1)$, then the possible tableaux of shape λ are the following:

$$t = \begin{array}{|c|c|}\hline 1 & 2 \\ \hline 3 \\ \cline{1-1} \end{array}, \quad \begin{array}{|c|c|}\hline 1 & 3 \\ \hline 2 \\ \cline{1-1} \end{array}, \quad \begin{array}{|c|c|}\hline 2 & 1 \\ \hline 3 \\ \cline{1-1} \end{array}, \quad \begin{array}{|c|c|}\hline 2 & 3 \\ \hline 1 \\ \cline{1-1} \end{array}, \quad \begin{array}{|c|c|}\hline 3 & 1 \\ \hline 2 \\ \cline{1-1} \end{array}, \quad \begin{array}{|c|c|}\hline 3 & 2 \\ \hline 1 \\ \cline{1-1} \end{array}.$$

Definition 0.6 Two λ-tableaux t_1 and t_2 are *row equivalent*, $t_1 \sim t_2$, if corresponding rows of the two tableaux contain the same elements. A *tabloid of shape* λ, or λ-tabloid, is then

$$\{t\} = \{t_1 \mid t_1 \sim t\}.$$

where the shape of t is λ.

EXAMPLE. If we have the λ-tableau

$$t = \begin{array}{|c|c|}\hline 2 & 3 \\ \hline 1 \\ \cline{1-1} \end{array},$$

then the corresponding λ-tabloid is

$$\{t\} = \{ \begin{array}{|c|c|}\hline 2 & 3 \\ \hline 1 \\ \cline{1-1} \end{array}, \quad \begin{array}{|c|c|}\hline 3 & 2 \\ \hline 1 \\ \cline{1-1} \end{array} \}.$$

An element $\pi \in S_n$ acts on a tableau t by sending each number in t through π individually. For example,

$$(132) \begin{array}{|c|c|}\hline 2 & 3 \\ \hline 1 \\ \cline{1-1} \end{array} = \begin{array}{|c|c|}\hline 1 & 2 \\ \hline 3 \\ \cline{1-1} \end{array}.$$

This induces an action on a tabloid $\{t\}$ by letting

$$\pi\{t\} = \{\pi t\}.$$

Definition 0.7 Suppose $\lambda \vdash n$. Let

$$M^\lambda = \mathbb{C}\{\ \mathbf{\{}t_1\mathbf{\}}\ ,\ldots,\ \mathbf{\{}t_k\mathbf{\}}\ \},$$

where $\{t_1\}, \ldots, \{t_k\}$ is a complete list of λ-tabloids. Then M^λ is called the *permutation module corresponding to* λ. S_n acts on M^λ by permuting $\{t_1\}, \ldots, \{t_k\}$.

Definition 0.8 Suppose that the tableau t has rows $R_1, R_2, \ldots, R_l$ and columns $C_1, C_2, \ldots, C_k$. Then

$$R_t = \mathcal{S}_{R_1} \times \mathcal{S}_{R_2} \times \cdots \times \mathcal{S}_{R_l}$$

and

$$C_t = \mathcal{S}_{C_1} \times \mathcal{S}_{C_2} \times \cdots \times \mathcal{S}_{C_k}$$

are the *row-stabilizer* and *column-stabilizer* of t, respectively.

EXAMPLE. If we have the λ-tableau

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

then

$$R_t = \mathcal{S}_{\{1,3,4,6\}} \times \mathcal{S}_{\{2,5\}}$$

and

$$C_t = \mathcal{S}_{\{2,3\}} \times \mathcal{S}_{\{5,6\}} \times \mathcal{S}_{\{4\}} \times \mathcal{S}_{\{1\}}.$$

We know that every $\pi \in S_n$ can be written as a product of transposition τ_i, $\pi = \tau_1\tau_2\cdots\tau_k$. For the following definition, we will need to make use of

$$\kappa_t := C_t^- = \sum_{\pi \in C_t} sgn(\pi)\pi,$$

where $sgn(\pi) = (-1)^k$. Since the parity (mod 2) of the number of transpositions needed to write π is well-defined, it follows that $(-1)^k$ is also well-defined.

Definition 0.9 If t is a tableau, then the associated *polytabloid* is $\mathbf{e}_t = \kappa_t\mathbf{\{t\}}$.

EXAMPLE. Again, if

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

then

$$\kappa_t = 1 - (23) - (56) + (23)(56).$$

So

$$\begin{aligned}
\mathbf{e}_t &= [1-(23)-(56)+(23)(56)]\{t\} \\
&= \{t\}-(23)\{t\}-(56)\{t\}+(23)(56)\{t\} \\
&= \{t\}-\{(23)t\}-\{(56)t\}+\{(23)(56)t\} \\
&= \{\begin{array}{|c|c|c|c|}\hline 3&6&4&1\\ \hline 2&5\\ \cline{1-2}\end{array}\} - \{(23)\begin{array}{|c|c|c|c|}\hline 3&6&4&1\\ \hline 2&5\\ \cline{1-2}\end{array}\} - \{(56)\begin{array}{|c|c|c|c|}\hline 3&6&4&1\\ \hline 2&5\\ \cline{1-2}\end{array}\} + \{(23)(56)\begin{array}{|c|c|c|c|}\hline 3&6&4&1\\ \hline 2&5\\ \cline{1-2}\end{array}\} \\
&= \{\begin{array}{|c|c|c|c|}\hline 3&6&4&1\\ \hline 2&5\\ \cline{1-2}\end{array}\} - \{\begin{array}{|c|c|c|c|}\hline 2&6&4&1\\ \hline 3&5\\ \cline{1-2}\end{array}\} - \{\begin{array}{|c|c|c|c|}\hline 3&5&4&1\\ \hline 2&6\\ \cline{1-2}\end{array}\} + \{\begin{array}{|c|c|c|c|}\hline 2&5&4&1\\ \hline 3&6\\ \cline{1-2}\end{array}\}.
\end{aligned}$$

Definition 0.10 For any partition λ, the corresponding *Specht module*, S^λ, is the submodule of M^λ spanned by the polytabloids $\mathbf{e}_t$, where t is of shape λ.

These S^λ constitute a full set of irreducible $\mathcal{S}_n$-modules if $char(F)=0$. If $char(F)=p$, one still defines Specht modules, but they are usually not irreducible. We will be studying what conditions guarantee that S^λ over F_p is irreducible.

In addition to looking at these Specht modules, we also want to focus on the diagram of λ (or its Young tableau), and what we can deduce from various properties about lengths of rows and columns in the tableau.

In the most intuitive terms, the *hook* of a node $v=(i,j)$ (the box in the ith row and the jth column) in the diagram of λ is the node v together with the set of nodes to the right of v in row i, and the set of nodes below v in column j. And thus the *hooklength* of v is just the number of such nodes. A *rim hook* shall be defined as being obtained by projecting a regular hook along diagonals onto the boundary of the diagram of λ.

EXAMPLE. If we have a λ-tableau with $v=(1,1)$

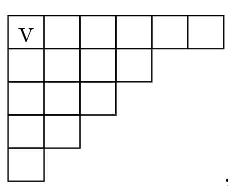

,

then we get the corresponding hook of v (denoted by the set of nodes with dots inside)

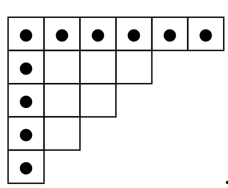

,

and corresponding rim hook of v (denoted by the set of nodes with crosses inside)

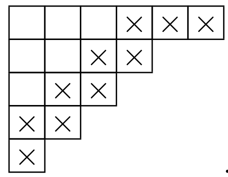

.

We should be able to see that a diagram of a partition λ is completely determined by its first-column hooklengths h_i. The method we use to construct λ from these hooklengths is as follows:

$$\lambda_k = h_k, \lambda_{k-1} = h_{k-1} - 1, \ldots, \lambda_1 = h_1 - k + 1,$$

where λ_i is the length of the ith row of the diagram of λ. Notice that in doing so, we get a partition λ from each sequence of strictly decreasing nonnegative integers

$$\beta_1 > \beta_2 > \cdots > \beta_r$$

by letting

$$\lambda_i := \beta_i + i - r, 1 \leq i \leq r.$$

Such β-numbers can be conveniently recorded on an abacus. Imagine an abacus lying on a table with runners going north-south. We assume there are p runners, called the 0th runner, 1st runner, ..., (p-1)th runner, from left to right. The possible bead positions are determined by assuming that all the beads are initially at the top and that we move beads only through one bead width at a time. Label the bead positions as below:

$$\begin{array}{ccccc} 0 & 1 & \cdots & p-2 & p-1 \\ p & p+1 & \cdots & 2p-2 & 2p-1 \\ . & . & & . & . \\ . & . & & . & . \\ . & . & & . & . \end{array}$$

A bead configuration is associated with a set of β-numbers (and hence a partition) by letting the actual bead positions determine the β-numbers. What we aim to show is that there is a direct correlation between what happens when we move beads on the abacus and what happens to the diagram of the corresponding partition. In particular, we want to see that the removal of a rim p-hook from a partition occurs if and only if a bead in the corresponding abacus is moved up one row on the same string, where a *p-hook* is simply a hook with hooklength a prime p.

Theorem 0.12 *The removal of a rim p-hook from a partition occurs if and only if a bead in the corresponding abacus is moved up one row on the same string.*

PROOF. Assume that we have removed a rim p-hook from a partition. Let α_n denote the length of the nth row of the partition. Call α_i the length of topmost row involved in the rim p-hook that was removed. Then for all rows j involved in the rim p-hook except the bottom-most row k,

$$\alpha_j \mapsto \alpha_{j+1} - 1,$$

since row j is necessarily shortened to one less than the length of row $(j+1)$ by definition of a rim p-hook. We know by definition the following:

$$\alpha_j = \beta_j + j - r,$$

where β_n denotes the position of the nth bead of our abacus. So

$$\begin{aligned} \alpha_{j+1} - 1 &= \beta_{j+1} + j + 1 - r - 1 \\ &= \beta_{j+1} + j - r. \end{aligned}$$

So we can determine to where we map each β_j:

$$\begin{aligned} \beta_j &= \alpha_j - j + r \\ &\mapsto \alpha_{j+1} - 1 - j + r \\ &= \beta_{j+1} + j - r - j + r \\ &= \beta_{j+1}. \end{aligned}$$

Now we need to determine to where we map β_k. So first we will determine to where we map α_k. From our map, we can see that for each row j, α_j is shortened by $\alpha_j - (\alpha_{j+1} - 1)$, which equals $\alpha_j - \alpha_{j+1} + 1$. So the total number of nodes removed by the non-kth rows of the p-hook is

$$\begin{aligned} \sum_{j=i}^{k-1} (\alpha_j - \alpha_{j+1} + 1) &= (\alpha_i - \alpha_{i+1} + 1) + \cdots + (\alpha_j - \alpha_{j+1} + 1) + \cdots + (\alpha_{k-1} - \alpha_k + 1) \\ &= \alpha_i - \alpha_k + 1 \cdot (k - 1 - i + 1) \\ &= \alpha_i - \alpha_k + k - i. \end{aligned}$$

Since a p-hook removes p total nodes, clearly α_k must be shortened by $p - (\alpha_i - \alpha_k + k - i)$. Thus,

$$\begin{aligned} \alpha_k &\mapsto \alpha_k - (p - [\alpha_i - \alpha_k + k - i]) \\ &= \alpha_k - p + \alpha_i - \alpha_k + k - i \\ &= \alpha_i + (k - i) - p. \end{aligned}$$

Using this fact we get

$$\begin{aligned} \beta_k &= \alpha_k - k + r \\ &\mapsto \alpha_i + k - i - p - k + r \\ &= \alpha_i - i + r - p \\ &= \beta_i - p. \end{aligned}$$

Since we have that $\beta_j \mapsto \beta_{j+1}$ and that $\beta_k \mapsto \beta_i - p$, this is equivalent to having moved the ith bead on our abacus up one row on the same string.

To prove the converse, assume that we have moved the ith bead on our abacus up one row on the same string. Then we have that $\beta_j \mapsto \beta_{j+1}$ and that $\beta_k \mapsto \beta_i - p$. So

$$\begin{aligned} \alpha_j &= \beta_j + j - r \\ &\mapsto \beta_{j+1} + j - r \\ &= \alpha_{j+1} - (j+1) + j - r \\ &= \alpha_{j+1} - 1. \end{aligned}$$

And

$$\begin{aligned} \alpha_k &= \beta_k + k - r \\ &\mapsto \beta_i - p + k - r \\ &= \alpha_i - i + r - p + k - r \\ &= \alpha_i + (k - i) - p. \end{aligned}$$

So we have removed a rim hook. To show that we have removed a rim p-hook, we again note that the number of nodes removed by the non-kth rows of the rim hook is

$$\sum_{j=i}^{k-1} (\alpha_j - \alpha_{j+1} + 1),$$

and the number of nodes removed by the kth row is

$$p - (\alpha_i - \alpha_k + k - i).$$

So the total number of nodes removed by our rim hook is the sum of these two quantities:

$$\begin{aligned} \sum_{j=i}^{k-1} (\alpha_j - \alpha_{j+1} + 1) + p - (\alpha_i - \alpha_k + k - i) &= \alpha_i - \alpha_k + k - i + p - (\alpha_i - \alpha_k + k - i) \\ &= \alpha_i - \alpha_k + k - i + p - \alpha_i + \alpha_k - k + i \\ &= p. \end{aligned}$$

Since we have removed p nodes with our rim hook, we must have removed a rim p-hook. □

Definition 0.13 A diagram of a partition λ is called a p-*core* if it does not contain any p-hooks.

Corollary 0.14 *A partition is a p-core if and only if no bead in the corresponding abacus can be moved up one row on the same string.*

PROOF. Assume a partition is a p-core. Then by definition the partition contains no p-hooks, so we cannot remove one. Since we cannot remove a p-hook, we cannot move a bead in the corresponding abacus up one row on the same string by the preceding theorem. So no bead in the corresponding abacus can be moved up one row on the same string.

Conversely, assume no bead on our abacus can be moved up one row on the same string. So we cannot remove a p-hook from our partition by the preceding theorem. So our partition is a p-core. □

Definition 0.15 A node (a, b) is *p-isolated* if there exists a node (a, x) in the same row as (a, b) and a node (y, b) in the same column as (a, b) such that the p-part (greatest power of p dividing the node) of (a, x) and the p-part of (y, b) are different than the p-part of (a, b).

Definition 0.16 A partition λ is *p-irreducible* if λ has no p-isolated nodes.

The *p-irreducible* partitions are the ones which will most help us in the representation theory of the symmetric group. We give a theorem here by Fayers and part of one direction of the proof, because it is paramount in our understanding of these partitions. In short, it says that a Specht module S^λ over F_p is irreducible if and only if the corresponding partition λ is *p-irreducible.*

The proposition below characterizes p-irreducible partitions in terms of their abaci.

Proposition 0.17 *The following are equivalent:*
(1) λ is a p-irreducible partition.
(2) There exist some i and j such that:
(a) $\lambda(k) = \emptyset$ whenever $i \neq k \neq j$,
(b) if position $i + pa$ on runner i is unoccupied, then any position $b > i + ea$ not on runner i is unoccupied,
(c) if position $j + pc$ on runner j is occupied, then any position $d < j + ec$ not on runner j is occupied,
(d) $\lambda(i)$ is a p-regular p-irreducible partition,
(e) $\lambda(j)$ is a conjugate p-restricted p-irreducible partition.

PROOF. We want to prove that (2) implies (1). So suppose that λ has an abacus configuration as described in (2), and suppose that $h_\lambda(a, c)$ is divisible by p, say $h_\lambda(a, c) = ps$. We *claim* that this means that there is an unoccupied space exactly s spaces above the bead corresponding to the beta-number β_a on the same runner. Hence this bead must lie either on runner i or runner j. We shall suppose that it lies on runner i (the case where it lies on runner j may be addressed by replacing λ with its conjugate). We claim that, for $b = 1, \ldots, \lambda'_c$ we have

$$v_p(h_\lambda(b, c)) = v_p(h_\lambda(a, c)),$$

where v_p is the p-part of the node. Write $d = \lambda'_c$. Since there is a node (a, c) it is clear that $d \geq a$. Since there is an unoccupied space exactly s spaces above the bead corresponding to β_a, this space is in position $\beta_a - ps$ on runner i. Then this is our $i + pa$ referred to by (2b). So any position $b > i + pa = \beta_a - ps$ not on runner i is unoccupied. We want to show that the beads corresponding to $\beta_1, \ldots, \beta_d$ all lie on runner i, so it suffices to show that $\beta_1, \ldots, \beta_d$ are all greater than $\beta_a - ps$. Since our β numbers are in strictly decreasing order, it suffices to show that $\beta_d > \beta_a - ps$. Assume that

$$\beta_d \leq \beta_a - ps.$$

This implies

$$\lambda_d - d + r \leq \lambda_a - a + r - ps$$
$$\Rightarrow \lambda_d - d \leq \lambda_a - a - ps$$
$$\Rightarrow \lambda_d + a + ps \leq \lambda_a + d.$$

Since we know that $h_\lambda(a,c) = ps$, we can explicitly determine $\lambda_a + d$:

$$\lambda_a + d = c - 1 + a - 1 + ps = c - 2 + a + ps.$$

So now

$$\lambda_d + a + ps \leq \lambda_a + d$$
$$\Rightarrow \lambda_d + a + ps \leq c - 2 + a + ps$$
$$\Rightarrow \lambda_d \leq c - 2.$$

But

$$\lambda_d = \lambda_{\lambda'_c} \geq c.$$

So we have

$$c \leq \lambda_d \leq c - 2,$$

which is clearly a contradiction. So we must have $\beta_d > \beta_a - ps$. So $\beta_1, \ldots, \beta_d$ are all greater than $\beta_a - ps = i + pa$, which implies that $\beta_1, \ldots, \beta_d$ all lie on runner i by condition (2b). So if we let M be the number of unoccupied spaces less than β_d on the abacus, we see simply that $\beta_d = M + r - d$, and so we have

$$\begin{aligned} \lambda_d &= M, \\ \lambda_{d-1} &= M + p(\tau_{d-1} - \tau_d) + p - 1, \\ &\vdots \\ \lambda_1 &= M + p(\tau_1 - \tau_d) + (d-1)(p-1). \end{aligned}$$

Put $y = \tau_a - s + d - a + 1$; then we *claim* that, for $x = 1, \ldots, d$,

$$h_\lambda(x,c) = eh_\tau(x,y),$$

this will then be sufficient, since τ is a (p,p)-Carter partition, so we have

$$v_{e,p}(h_\lambda(x,c)) = 1 + v_p(h_\tau(x,y)) = 1 + v_p(h_\tau(a,y)) = v_{e,p}(h_\lambda(a,c)).$$

First we claim that $\tau'_y = d$; this follows easily from the fact that $\lambda'_c = d$. Verifying the above equality is then a formality. □

Using conclusions like those made in Fayers theorem, in the future, we look to learn more about the representations of the symmetric group. The following idea of generating functions is a central part of what we plan to do in the near future.

Definition 0.18 Given a sequence of complex numbers $a_n = a_0, a_1, a_2, \ldots$, the corresponding *generating function* is the power series

$$f(x) = \sum_{n\geq 0} a_n x^n.$$

The generating function for partitions

$$\sum_{n\geq 0} p(n)x^n = \prod_{i\geq 1} \frac{1}{1-x^i},$$

where $p(n)$ is the number of partitions of n, is known. One can see that this is the generating function for partitions by looking closely at the function:

$$\begin{aligned}
(1)\sum_{n\geq 0} p(n)x^n &= \prod_{i\geq 1} \frac{1}{1-x^i} \\
(2) \qquad\qquad &= \prod_{i\geq 1}(1 + x^i + x^{2i} + x^{3i} + \cdots) \\
(3) \qquad\qquad &= (1 + x + x^2 + x^3 + \cdots)(1 + x^2 + (x^2)^2 + \cdots)(1 + x^3 + (x^3)^2 + \cdots)\cdots \\
(4) \qquad\qquad &= 1x^0 + 1x^1 + 2x^2 + 3x^3 + 5x^4 + 7x^5 + \cdots.
\end{aligned}$$

We can see from line (3) how the coefficients of this function count the partitions of each n. When $i = 1$, the coefficients are counting the number of ways the number 1 appears in the partition of n; when $i = 2$, the coefficients are counting the number of ways the number 2 appears in the partition of n; and so on. For instance, the partitions of 3 are (3), (2,1), and (1,1,1). The (3) comes from the x^3 term when $i = 3$, the (2,1) comes from the x^2 term when $i = 2$ multiplied by the x term when $i = 1$, and the (1,1,1) comes from the x^3 term when $i = 1$.

There are known generating functions for many types of partitions, including the number of partitions of n in which each term is odd, even, a square, a prime, etc. What we aim to do is to determine a generating function for p-irreducible partitions, for this will lead to irreducible Specht modules and further to irreducible representations of S_n.

Before we do this, however, we need to determine more properties of partitions, including the generating function for p-cores. To help accomplish this, we look at a problem from the first chapter of I.G. MacDonald's *Symmetric Functions and Hall Polynomials*. It outlines the formation of such a generating function.

#8

(A)

Proposition 0.19 *Let λ, μ be partitions of length $\leq m$ such that $\lambda \supset \mu$, and such that $\lambda - \mu$ is a rim hook of length p. Let $\delta_m = (m-1, m-2, \ldots, 1, 0)$ and let $\xi = \lambda + \delta_m$, $\eta = \mu + \delta_m$. Then η is obtained from ξ by subtracting p from some part ξ_i of ξ and rearranging in descending order.*

PROOF. Let $\mu = (\mu_1, \ldots, \mu_k)$ and let $\lambda = (\mu_1 + \alpha_1, \ldots, \mu_k + \alpha_k, \alpha_{k+1}, \ldots, \alpha_{k+l})$. Now let j be the first row of the rim hook (i.e., α_j is the first nonzero α), and let h be the last row of the rim hook. We know that there are two cases: either $l > 0$ and thus the rim hook extends to the last row of λ, which implies that $k + l = h$; or $l = 0$, and thus the length of λ and the length of μ are equal. The following statements where $i \neq h$ hold true for both cases.

Consider i such that $i < j$. Then it is clear that

$$\eta_i = \xi_i.$$

Now consider i such that $j \leq i < h$. By definition

$$\begin{aligned} \alpha_i &= \lambda_i - \lambda_{i+1} + 1 \\ &= \mu_i + \alpha_i - \mu_{i+1} - \alpha_{i+1} + 1. \end{aligned}$$

This implies

$$\begin{aligned} \mu_i &= \mu_{i+1} + \alpha_{i+1} - 1 \\ &= \lambda_{i+1} - 1. \end{aligned}$$

So

$$\begin{aligned} \eta_i &= \mu_i + (m - i) \\ &= \lambda_{i+1} - 1 + (m - i) \\ &= \lambda_{i+1} + (m - i - 1) \\ &= \xi_{i+1}. \end{aligned}$$

Next consider i such that $i > h$. Again, in either case

$$\begin{aligned} \eta_i &= \mu_i + (m - i) \\ &= 0 + (m - i) \\ &= m - i \\ &= 0 + m - i \\ &= \alpha_i + (m - i) \\ &= \xi_i. \end{aligned}$$

Our final case is for the hth row. We can see that

$$\eta_h = \mu_h + (m - h).$$

However, we must now split into two cases.

CASE 1: $l = 0$. Then $\mu_h = \lambda_h - \alpha_h$. In this case, we can see that p is equal to the difference in hooklengths of $(j, 1)$ and $(h, 1) + \alpha_h$. In other words,

$$\begin{aligned} p &= \lambda_j - j + \lambda'_1 - 1 + 1 - (\lambda_h - h + \lambda'_1 - 1 + 1) + \alpha_h \\ &= \lambda_j - j - \lambda_h + h + \alpha_h. \end{aligned}$$

So

$$\begin{aligned} \eta_h &= \mu_h + (m - h) \\ &= \lambda_h - \alpha_h + (m - h) \\ &= \lambda_j - j - (\lambda_j - j - \lambda_h + h + \alpha_h) + m \\ &= \lambda_j - j - p + m \\ &= \lambda_j + (m - j) - p \\ &= \xi_j - p. \end{aligned}$$

CASE 2: $l > 0$. Then $\mu_h = 0$ and $k + l = h$. In this case, we can see that p is equal to the hooklength of $(j, 1)$. In other words,

$$\begin{aligned} p &= \lambda_j - j + \lambda'_1 - 1 + 1 \\ &= \lambda_j - j + k + l \\ &= \lambda_j - j + h. \end{aligned}$$

So

$$\begin{aligned} \eta_h &= \mu_h + (m - h) \\ &= 0 + (m - h) \\ &= m - h \\ &= \lambda_j - j - (\lambda_j - j + h) + m \\ &= \lambda_j - j - p + m \\ &= \lambda_j + (m - j) - p \\ &= \xi_j - p. \end{aligned}$$

So in either case, $\eta_h = \xi_j - p$. Thus η is obtained from ξ by subtracting p from ξ_j of ξ and rearranging in descending order. □

(B)

Definition 0.20 With the same notation as in #8(a), suppose that ξ has m_r parts ξ_i congruent to r modulo p, for each $r = 0, 1, \ldots, p-1$. These ξ_i may be written in the form $p\xi_k^{(r)} + r$ $(1 \leq k \leq m_r)$, where $\xi_1^{(r)} > \xi_2^{(r)} > \ldots > \xi_{m_r}^{(r)} \geq 0$. Let $\lambda_k^{(r)} = \xi_k^{(r)} - m_r + k$, so that $\lambda^{(r)} = (\lambda_1^{(r)}, \ldots, \lambda_{m_r}^{(r)})$ is a partition. The collection $\lambda^* = (\lambda^{(0)}, \lambda^{(1)}, \ldots, \lambda^{(p-1)})$ is called the *p-quotient* of the partition λ.

Proposition 0.21 *The preceding definition of the p-quotient of λ is equivalent to our previous one involving an abacus.*

PROOF. We first claim that the parts ξ_i from Definition 2 correspond exactly to our β-numbers in the abacus definition. We know that $\xi_i = \lambda_i + m - i$ by definition. Here, m is simply equal to our r from the abacus definition (i.e., the number of β-numbers). To see the abacus correspondence, note that if we plot the ξ_i on an abacus with p runners labeled $0, 1, \dots, p-1$ we get that if $\xi_i = p\xi_k^{(r)} + r$, then ξ_i is on the $\xi_k^{(r)}$th row of the rth runner (if we start labeling rows at 0). Also, from Definition 2 it is clear that plotting this way puts m_r beads on the rth runner. So the kth row of the rth partition in the p-quotient, which is given by $\lambda_k^{(r)} = \xi_k^{(r)} - m_r + k$, is the row of the kth bead on the rth runner minus the total number of beads on the rth runner plus the number of beads greater in value than the kth bead on the rth runner. In other words, this formula exactly counts the number of empty positions on the rth runner less than the kth bead on the rth runner. This is how we define the p-quotient in our abacus definition, so the two are equivalent. $\square$

Definition 0.22 The m numbers $ps + r$, where $0 \leq s \leq m_r - 1$ and $0 \leq r \leq p - 1$, are all distinct. Let us arrange them in descending order, say $\widetilde{\xi_1} > \cdots > \widetilde{\xi_m}$, and define a partition $\widetilde{\lambda}$ by $\widetilde{\lambda}_i = \widetilde{\xi}_i - m + i$ $(1 \leq i \leq m)$. This partition $\widetilde{\lambda}$ is called the *p-core* of λ. Both $\widetilde{\lambda}$ and λ^* (up to cyclic permutation) are independent of m, provided that $m \geq l(\lambda)$.

If $\lambda = \widetilde{\lambda}$ (i.e. if λ^* is empty), the partition λ is called a *p-core*. For example, the only 2-cores are the 'staircase' partitions $\delta_m = (m-1, m-2, \dots, 1)$.

Following G.D. James, we may conveniently visualize this construction in terms of an abacus. The runners of the abacus are the half-lines $x \geq 0$, $y = r$ in the plane $\mathbf{R}^2$, where $r = 0, 1, 2, \dots, p-1$, and λ is represented by the set of beads at the points with coordinates $(\xi_k^{(r)}, r)$ in the notation used above. The removal of a rim hook of length p from λ is recorded on the abacus by moving some bead one unit to the left on its runner, and hence the passage from λ to its p-core corresponds to moving all the beads on the abacus as far left as they will go.

Recall that we proved this in Theorem 0.12.

This arithmetical construction of the p-quotient and p-core is an analogue for partitions of the division algorithm for integers (to which it reduces if the partition has only one part).

(C)

The p-core of a partition λ may be obtained graphically as follows. Remove a rim hook of length p from the diagram of λ in such a way that what remains is the diagram of a partition, and continue removing rim hooks of length p in this way as long as possible. What remains at the end of this process is the p-core of $\widetilde{\lambda}$ of λ, and it is independent of the sequence of rim hooks removed. For by (a) above, the removal of a rim hook of length p from λ corresponds to subtracting p from some part of ξ and then rearranging the resulting sequence in descending order; the only restriction is that the resulting set of numbers should be all distinct and non-negative.

Recall that we proved this exact statement in Fayers's Proposition.

(D)

The p-quotient of λ can also be read off from the diagram of λ, as follows. For s, $t = 0, 1, \ldots, p-1$ let

$$R_s = \{(i,j) \in \lambda : \lambda_i - i \equiv s(modp)\},$$

$$C_t = \{(i,j) \in \lambda : j - \lambda'_j \equiv t(modp)\},$$

so that R_s consists of the rows of λ whose right-hand node has content congruent to s modulo p, and likewise for C_t. Recall that the content of a node (i,j) is equal to $j-i$, so since $j = \lambda_i$ and $i = \lambda'_j$ in the right-hand node, the preceding statement is true by definition. If now $(i,j) \in R_s \cap C_t$, the hook-length at (i,j) is

$$h(i,j) = \lambda_i + \lambda'_j - i - j + 1 \equiv s - t + 1(modp)$$

and therefore p divides $h(i,j)$ if and only if $t \equiv s + 1(modp)$.

On the other hand, if $\xi_i = p\xi_k^{(r)} + r$ as in part (b), then the hooklengths of λ in the ith row are the elements of the sequence $(1, 2, \ldots, \xi_i)$ after deletion of $\xi_i - \xi_{i+1}, \ldots, \xi_i - \xi_m$. In other words, there are no nodes of λ in the ith row with hooklength equal to $\xi_i - \xi_k$, where $i + 1 \leq k \leq m$. In other words we have no hooklengths equal to $\lambda_i - \lambda_k + (k-i)$. To see this, note the following:

$$\begin{aligned} h(i, \lambda_k) &= \lambda_i + \lambda'_k - i - \lambda_k + 1 \\ &\geq \lambda_i + k - i - \lambda_k + 1 \\ &= \lambda_i - \lambda_k + (k-i) + 1 \\ &> \lambda_i - \lambda_k + (k-i) \\ &= \lambda_i + k - i - \lambda_k \\ &> \lambda_i + (\lambda_k + 1)' - i - \lambda_k - 1 + 1 \\ &= h(i, \lambda_k + 1). \end{aligned}$$

This eliminates $m-i$ numbers, leaving $\xi_i - (m-i) = \lambda_i + (m-i) - (m-i) = \lambda_i$ numbers - exactly the number of nodes in the ith row. And hence those divisible by p are the elements of the sequence $(p, 2p, \ldots, p\xi_k^{(r)})$ after deletion of $p(\xi_k^{(r)} - \xi_{k+1}^{(r)}), \ldots, p(\xi_k^{(r)} - \xi_{m_r}^{(r)})$. They are therefore p times the hooklengths in the kth row of $\lambda^{(r)}$, and in particular there are $\lambda_k^{(r)}$ of them. In other words, there are no nodes of $\lambda^{(r)}$ in the kth row with hooklength equal to $\xi_k^{(r)} - \xi_j^{(r)}$, where $i + 1 \leq j \leq m_r$. In other words, we have no hooklengths equal to

$\xi_k^{(r)} - \xi_j^{(r)} + (j-k)$. To see this, again note the following:

$$\begin{aligned} h(k, \xi_j^{(r)}) &= \xi_k^{(r)} + \xi_j^{(r)\prime} - k - \xi_j^{(r)} + 1 \\ &\geq \xi_k^{(r)} + j - k - \xi_j^{(r)} + 1 \\ &= \xi_k^{(r)} - \xi_j^{(r)} + (j-k) + 1 \\ &> \xi_k^{(r)} - \xi_j^{(r)} + (j-k) \\ &= \xi_k^{(r)} + j - k - \xi_k^{(r)} \\ &> \xi_k^{(r)} + (\xi_j^{(r)} + 1)' - k - \xi_j^{(r)} - 1 + 1 \\ &= h(k, \xi_j^{(r)} + 1). \end{aligned}$$

This eliminates $m_r - k$ numbers, leaving $\xi_k^{(r)} - (m_r - k) = \xi_k^{(r)} - m_r + k = \lambda_k^{(r)}$ numbers - exactly the number of nodes in the k row of $\lambda^{(r)}$.

Since the ith row of λ corresponds to the kth row of $\lambda^{(r)}$, we have that if the ith row of λ is in R_s, then the kth row of $\lambda^{(r)}$ is in R_s. So we get the following:

$$\begin{aligned} s &\equiv \lambda_i - i(modp) \\ &\equiv \xi_i - m + i - i(modp) \\ &\equiv \xi_i - m(modp) \\ &\equiv p\xi_k^{(r)} + r - m(modp) \\ &\equiv r - m(modp). \end{aligned}$$

If we take this combined with the aforementioned condition for p dividing $h(i,j)$ if and only if $t \equiv s+1 \pmod p$, then it follows that each $\lambda^{(r)}$ is embedded in λ as $R_s \cap C_{s+1}$, where $s \equiv r - m \pmod p$, and that the hooklengths in $\lambda^{(r)}$ are those of the corresponding nodes in $R_s \cap C_{s+1}$, divided by p. In particular, if m is a multiple of p (which we may assume without loss of generality) then $\lambda^{(r)} = \lambda \cap R_r \cap C_{r+1}$ for each r (where $C_p = C_0$).

(E)

From (c) and (d) it follows that the p-core (respectively, p-quotient) of the conjugate partition λ' is the conjugate of the p-core (respectively, p-quotient) of λ.

(F)

For any two partitions λ, μ we shall write

$$\lambda \sim_p \mu$$

to mean that λ and μ have the same p-core.

Proposition 0.23 *As above, let $\xi = \lambda + \delta_m$, $\eta = \mu + \delta_m$, where $m \geq max(l(\lambda), l(\mu))$. Then it follows from (a) and (b) that $\lambda \sim_p \mu$ if and only if $\eta \equiv w\xi$ (mod p) for some permutation $w \in S_m$.*

PROOF. Assume $\lambda \sim_p \mu$. Then λ and μ have the same p-core. By the way we defined p-core in (b), this means that for every r, m_r is equal for ξ and η. So it seems that $\eta \equiv \xi$ (mod p). Now conversely assume that $\eta \equiv w\xi$ (mod p) for some permutation $w \in S_m$. Then the result clearly follows. □

Also, from (e) above it follows that $\lambda \sim_p \mu$ if and only if $\lambda' \sim_p \mu'$.

(G)

From the definitions in (b) it follows that a partition λ is uniquely determined by its p-core $\widetilde{\lambda}$ and its p-quotient λ^*. Since $|\lambda| = |\widetilde{\lambda}| + p|\lambda^*|$, the generating function for partitions with a given p-core $\widetilde{\lambda}$ is

$$\sum_{\widetilde{\mu}=\widetilde{\lambda}} t^{|\mu|} = t^{|\widetilde{\lambda}|} P(t^p)^p$$

where $P(t) = \prod_{n \geq 1}(1-t^n)^{-1}$ is the partition generating function. Note that the p on the inside of the parentheses accounts for the fact that each node in the p-quotient accounts for p nodes in the actual partition, and that the p on the outside of the parentheses accounts for the fact that there are p p-quotients. (This can again be visualized as the p runners on the abacus.)

Hence the generating function for p-cores is

$$\begin{aligned} \sum t^{|\widetilde{\lambda}|} &= P(t)/P(t^p)^p \\ &= \prod_{n \geq 1} \frac{(1-t^{np})^p}{1-t^n}. \end{aligned}$$

In particular, when $p = 2$, we obtain the identity

$$\sum_{m \geq 1} t^{m(m-1)/2} = \prod_{n \geq 1} \frac{1-t^{2n}}{1-t^{2n-1}}.$$

Learning what we have from this problem in the first chapter of MacDonald's *Symmetric Functions*, we can now proceed toward a major goal: that is, to find the generating function for p-irreducible partitions.

Doing this will require a number of smaller steps beforehand. First, we want to determine the generating function for the number of p-irreducible partitions with a given p-core, where the only nontrivial runner[1] is the ith runner, as described in the Fayers paper. Furthermore,

[1]This is tantamount to the only nontrivial p-quotient.

we want to determine the generating function for the number of p-irreducible partitions with a given core, where the only nontrivial runner[2] is the jth runner. From these two functions, I think it will be possible to determine the generating function for the total number of p-irreducible partitions with a given p-core.

SOURCES:

Fayers, Matthew. "Irreducible Specht modules for Hecke algebras of type A." *Advances in Mathematics*. 193 (2005): 438-452.

James, Gordon and Kerber, Adalbert. *The Representation Theory of the Symmetric Group*. Reading: Addison-Wesley Publishing Company Inc., 1981.

MacDonald, I.G. *Symmetric Functions and Hall Polynomials*. New York: Oxford University Press, 1995.

Sagan, Bruce E. *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*. New York: Springer-Verlag, 2001.

[2]This is tantamount to the only nontrivial p-quotient.