Math. 126 Lecture 15

The induction ring of the symmetric groups.

November 13, 2002

Contents

1 The Grothendieck group of a group.

Let G be a group and U, V, W etc. be G-modules, that is vector spaces together with a representation of G on them. Consider the space of all integer combinations of the V (thought of as abstract symbols, for example) so this is some huge abelian group (i.e. a $\mathbb{Z}$-module). Quotient this out by the subspace spanned by expessions

$$V - U - W$$

whenever there is an exact sequence

$$0 \to U \to V \to W \to 0$$

of G-modules. The image of V in this quotient space will be denoted by $[V]$.

In this course we have only considered finite groups G so far, and so will restrict ourselves to finte dimensional vector spaces over the complex numbers. A more intuitive way of saying what we are is that we are identifying two equivalent representations - so $[V]$ stands for the equivalence class of all representations equivalent to V, and if

$$V = U \oplus W$$

then $[V] = [U] + [W]$. (This is on account of Mashke's theorem which says that any invariant subspace has an invariant complement.) The new think that we are introducing is being able to subtract representations, of $-[V]$ is legitimate gadget. The space of such formal combinations of equivalence classes of representations (modulo the above equivalence relation) is called the **Grothendieck group** of G. An element of the Grothendieck group of G is called a **virtual representation**.

In this lecture we are interested in the case where G is a symmetric group, say $G = S_k$, in which case the irreducible representations of G are the Specht modules F^λ as λ ranges over all Young diagrams with k boxes. We will write

$$[\lambda] \quad \text{insteadof} \quad [F^\lambda].$$

2 The induction ring of the symmetric groups.

The direct product $S_p \times S_q$ can be considered as the subgroup of S_{p+q} consisting of those permutations which permute the first p and last q letters among themselves. So if V is an S_p module and W is an S_q module, then $V \otimes W$ is an $S_p \times S_q$ module, and so we can consider the induced module

$$(V \otimes W) \uparrow \quad \text{of} \quad S_{p+q}.$$

We will define a multiplication

$$[V][W] = [(V \otimes W) \uparrow].$$

It is clear that this multiplication is well defined (does not depend on the equivalence class of V) and is distributive over addition.

So if we let $\mathcal{R}^k$ denote the Grothendieck group of S_k, we have defined a biinear map

$$\mathcal{R}^p \times \mathcal{R}^q \to \mathcal{R}^{p+q}.$$

We thus have made

$$\mathcal{R} := \bigoplus_k \mathcal{R}^k$$

into a graded ring.

3 Isomorphism with the ring of symmetric functions.

The obvious basis to take of $\mathcal{R}^k$ is the set of $[\lambda]$ as λ ranges over all Young diagrams with k boxes. But remember how we constructed F^λ. We first constructed M^λ which was the space of functions on the tabloids of shape λ, and F^λ occurs once in M^λ, and all other irreducibles in M^λ correspond to "earlier" diagrams. This means that the $[M^\lambda]$ form a basis of $\mathcal{R}^k$. If $\lambda = (\lambda_1, \dots, \lambda_r)$ then by its very definition, M^λ is the induced representation from the trivial representation of $S_{\lambda_1} \times S_{\lambda_2} \cdots \times S_{\lambda_r}$. The trivial representation of S_ℓ is just the representation corresponding to the row (ℓ). In other words, in $\mathcal{R}$ we have

$$[M^\lambda] = [(\lambda_1)][(\lambda_2)] \cdots [(\lambda_r)]$$

and the $[M^\lambda]$ form a basis of $\mathcal{R}$.

In Λ, the Schur function corresponding to a row (ℓ) is the complete symmetric function h_ℓ. We know that these generate Λ and their products

$$h_\lambda = h_{\lambda_1} h_{\lambda_2} \cdots h_{\lambda_r}$$

form a basis of Λ. So the map ϕ given by

$$\phi(h_\lambda) = [M^\lambda]$$

(and extending linearly) is an isomorphism of algebras.

The rest of this lecture will be devoted to proving the following fact:

$$\phi(s_\lambda) = [\lambda]. \tag{1}$$

4 Warmup.

Explicit computation comparing the power sum basis and the symmetric monomial basis shows that

$$\begin{array}{rcccc} p_{(3)} & = & x_1^3 + x_2^3 + \cdots & = & m_{(3)} \\ p_{(2,1)} & = & (x_1^2 + x_2^2 + \cdots)(x_1 + x_2 + \cdots) & = & m_{(3)} + m_{(2,1)} \\ p_{(1,1,1)} & = & (x_1 + x_2 + \cdots)^3 & = & m_{(3)} + 3m_{(2,1)} + 6m_{(1,1,1)}. \end{array}$$

The diagram (3), which is the maximal diagram among diagrams with three boxes gives rise to the module $M^{(3)}$ which consists of functions of a single point. The value of the character of the representation of S_3 on a three cycle (or any element) of S_3 is one, and this is the coefficient of $m_{(3)}$ on the right hand side of the first equation. In fact, let us make a little character table for the characters of the M^λ (as λ ranges over diagrams with three boxes) on the three conjugacy classes of S_3 remembering that the character is given by the number of fixed points of the group element acting on the space of tabloids:

$\lambda \quad \backslash \mu$	e	$3(ab)$	$2(abc)$
(3)	1	1	1
$(2,1)$	3	1	0
$(1,1,1)$	6	0	0

In this case,

Proposition 4.1 *If we expand the power sum basis in terms of the monomial basis,*

$$p_\mu = \sum \xi_\mu^\lambda m_\lambda \tag{2}$$

then the coefficient ξ_μ^λ is the value of the character of M^λ on the conjugacy class whose cycle structure is given by μ.

Proof in general. By definition ξ_μ^λ is the coefficient of x^λ in the expansion of

$$(x_1^{\mu_1} + \cdots)(x_2^{\mu} + \cdots) \cdots (x_1^{\mu_k} + \cdots) \cdots .$$

This coefficient will be the number of ways of choosing disjoint subsets of the μ_i so that the sum of the μ_i in the first subset equals λ_1, the sum of the μ_j in the second subset equals λ_2 etc.

Now consider the number of tabloids of type λ fixed by a $w \in S_n$ whose cycle structure is given by μ. A tabloid $\{t\}$ is fixed by w if and only if each cycle of w lies in the same row of t. So we must distribute the cycles of length μ_i in such a way that they fill up the rows (of length λ_ℓ) of λ, and this is the same combinatorial problem as in the expansion problem. QED

5 A key formula.

What we are going to need for the proof that $\phi(s_\lambda) = [\lambda]$ is not the expansion of the p in terms of the m but rather the expansion of the h in terms of the p. Here is how this is going to work: The space Λ^d can be considered as the space of integer combinations of characters of S_d which are just certain kinds of functions on S_d. There is a scalar product on this space functions, and the irreducible characters form an orthonormal system (an orthonormal basis of the space of central functions). If we believe that $\phi(s_\lambda) = [\lambda]$, then we might as well put a "scalar product" on Λ by declaring that $\Lambda^k \perp \Lambda^\ell$ if $k \neq \ell$ that the s_λ (as λ ranges over all partitions with d boxes) forms an orthonormal basis of Λ^d. It will turn out that relative to this scalar product we have

$$(h_\lambda, m_\mu) = \delta_{\lambda\mu} \tag{3}$$

and

$$(p_\lambda, p_\mu) = z(\lambda)\delta_{\lambda\mu} \tag{4}$$

where $z(\lambda)$ is our old friend, the denominator in the formula for the number of elements in the conjugacy class with cycle structure λ:

$$z(\lambda) = \prod_r r^{m_r} m_r!$$

where m_r is the number of times that r occurs in λ. If we now take the scalar product of (2) with h_λ we obtain

$$(h_\lambda, p_\mu) = \xi_\mu^\lambda$$

and hence

$$h_\lambda = \sum_\mu \frac{1}{z(\mu)} \xi^\lambda_\mu p_\mu. \qquad (5)$$

The is the key equation that we want. Of course, in order for it to even make sense, we have to be sure that our ring of coefficients contains the rational numbers.

Let us assume this key equation for the moment. This then says that the inverse ψ of ϕ sends $[M^\lambda]$ to $h_\lambda = \sum_\mu \frac{1}{z(\mu)} \xi^\lambda_\mu p\mu$ and since the $[M^\lambda]$ form a basis, that

$$\psi([V]) := \phi^{-1}([V]) = \sum_\mu \frac{1}{z(\mu)} \chi_V(C(\mu)) p_\mu \qquad (6)$$

where $C(\mu)$ denotes the conjugacy class consisting of those permutations whose cycle structure is given by μ. (The expression $\chi_V(C(\mu))$ then means the (constant) value that the function χ_V takes on the elements of this conjugacy class.)

Then

$$(\psi([V]), \psi([W]) = \sum_{\mu,\nu} \frac{1}{z(\mu)z(\nu)} \chi_V(C(\mu))\chi_W(C(\nu))(p_\mu, p_\nu) = \sum_\mu \frac{1}{z(\mu)} \chi_V(C(\mu))\chi_W(C(\nu))$$

by (4). But

$$\sum_\mu \frac{1}{z(\mu)} \chi_V(C(\mu))\chi_W(C(\nu)) = \frac{1}{n!} \sum_\mu \frac{n!}{z(\mu)} \chi_V(C(\mu))\chi_W(C(\nu)) = (\chi_V(C(\mu)), \chi_W(C(\nu)))_{S_d}.$$

The shows that the map ψ, and hence the map ϕ is an isometry. Now

$$h_\lambda = s_\lambda + \sum_{\mu \succ \lambda} K_{\mu,\lambda} s_\mu \quad \text{and} \quad [M^\lambda] = [\lambda] + \sum_{\mu \succ \lambda} k_{\mu\lambda}[[\mu]$$

which imply that there are integers $m_{\mu\lambda}$ such that

$$\phi(s_\lambda) = [\lambda] + \sum_{\mu \succ \lambda} m_{\mu,\lambda}[\mu].$$

Since we have proved that ϕ is an isometry, we have

$$1 = (s_\lambda, s_\lambda) = (\phi(s_\lambda), \phi(s_\lambda)) = 1 + \sum_\mu m^2_{\mu,\lambda}$$

so the $m_{\mu\lambda} = 0$ and we have the desired result:

$$\phi(s_\lambda) = [\lambda]. \qquad (7)$$

Of course, this hinges on our proving the key equation (5) and several earlier equations computing scalar products.

5.1 A special case of the key equation.

Consider the case where $\lambda = (d)$, a single row. In this case the representation is the trivial representation, and so $\xi_\mu^{(d)} = 1$ for all μ. We wish to prove

$$h_d(X_1, \ldots, X_n) = \sum_\mu \frac{1}{z(\mu)} p_\mu(X_1, \ldots, X_n) \tag{8}$$

in the ring of symmetric polynomials. We can prove this using generating functions:

$$\begin{aligned}
\sum_{d=0}^{\infty} h_d(X) t^d &= \prod_1^n \frac{1}{1 - X_i t} \\
&= \prod_1^n \exp(-\log(1 - X_i t)) \\
&= \prod_1^n \exp\left(\sum_{r=1}^{\infty} \frac{(X_i t)^r}{r}\right) \\
&= \exp\left(\sum_{r=1}^{\infty} \sum_{i=1}^{n} \frac{(X_i t)^r}{r}\right) \\
&= \exp\left(\sum_{r=1}^{\infty} \frac{p_r(X) t^r}{r}\right) \\
&= \prod_{r=1}^{\infty} \exp\left(\frac{p_r(X) t^r}{r}\right) \\
&= \prod_{r=1}^{\infty} \sum_{m_r=0}^{\infty} \frac{(p_r(X) t^r)^{m_r}}{m_r! r^{m_r}} \\
&= \sum_\mu \frac{1}{z(\mu)} t^{|\mu|} \quad \text{QED.}
\end{aligned}$$

5.2 Three expansions of the same of the same product.

Let $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_k$ be variables. Consider the product

$$\prod_{i=1}^{n} \prod_{j=1}^{k} \frac{1}{1 - X_i Y_j}.$$

If we first do the product over the X_i and collect coefficients of Y_j^d we see that this product is equal to

$$\prod_j \left(\sum h_d(X) Y_j^d\right).$$

Now collecting the coefficient of Y^λ and of the elements of its orbit under S_k we see that

$$\prod_{i=1}^{n}\prod_{j=1}^{k}\frac{1}{1-X_iY_j}=\sum_\lambda h_\lambda(X)m_\lambda(Y). \tag{9}$$

On the other hand, if we consider the X_iY_j as nk variables, we have

$$\prod_{i=1}^{n}\prod_{j=1}^{k}\frac{1}{1-X_iY_j}.=\sum_n h_n(X_1Y_1,X_1,Y_2,\dots,X_nY_k)=\sum_\mu\frac{1}{z(\mu})p_\mu(X_1Y_1,\dots,X_nY_k)$$

by (8). Now $p_r(X_1Y_1,\dots,X_nY_k)=p_r(X_1,\dots,X_n)p_r(Y_1,\dots,Y_k)$ and hence the same is true for p_λ. So we get our second expression for the same product:

$$\prod_{i=1}^{n}\prod_{j=1}^{k}\frac{1}{1-X_iY_j}=\sum_\mu\frac{1}{z(\mu)}p_\mu(X)p_\mu(Y). \tag{10}$$

The third formula for this product requires a different method of proof. We will follow Sagan and prove this by a variation of the Shensted bumping algorithm. Littlewood says that this formula goes back to Cauchy. I haven't checked this out but we will call it Cauchy's formula and use it before we prove it. The formula is

$$\prod_{i=1}^{n}\prod_{j=1}^{k}\frac{1}{1-X_iY_j}=\sum_\lambda s_\lambda(X)s_\lambda(Y). \tag{11}$$

5.3 Proof of the key formulas.

We have put a scalar product on Λ by declaring the s_λ to be orthonormal. Expand h and m in terms of the s:

$$h_\lambda=\sum a_{\lambda\mu}s_\mu,\qquad m_\lambda=\sum b_{\lambda\nu}s_\nu.$$

Substituting these expansions into the right hand side of (9) gives

$$\sum a_{\lambda\mu}b_{\lambda\nu}s_\mu s_\nu$$

which must equal the right hand side of (11) implying that

$$A^tB=I$$

where $A=(a_{\lambda\mu})$ and $B=(b_{\lambda\nu})$. This implies that $AB^t=I$ which says that

$$(h_\lambda,m_\tau)=(\sum_\mu a_{\lambda\mu}s_\mu,\sum_\nu b_{\tau\nu}s_\nu)=\sum a_{\lambda\mu}b_{tau,\nu}=\delta_{\mu,\nu}$$

which is (3).

If we write

$$p_\mu=\sum\chi^\lambda_\mu s_\lambda \tag{12}$$

Then comparing the right hand side of (10) with the right hand side of (11) establishes (4).

5.4 The Littlewood-Richardson rule for induced representations.

Using the isomorphism we can translate the Littlewood-Richardson rule from a theorem about multiplication of Schur functions into a rule about inducing from $S_k \times S_\ell$ to $S_{k+\ell}$.

5.5 The Frobenius formula for the character of $[\lambda]$.

From equations (12) and (4) we conclude that

$$s_\lambda = \sum_\mu \frac{1}{z(\mu)} \chi^\lambda_\mu p_\mu. \tag{13}$$

Apply the formula (6) to $[\lambda]$. We know that $\psi([\lambda]) = s_\lambda$. We conclude that that if χ^λ denotes the character of $[\lambda]$, then the value of χ^λ on the conjugacy class given by the cycle structure μ is χ^λ_μ, where χ^λ_μ is the coefficient occurring in (12).

We still must prove Cauchy's formula (11).

6 The Schensted-Knuth correspondence and Cauchy's formula.

6.1 Generalized permutations.

A **generalized permutation** is a two line array of positive integers whose columns are in lexicographic order with the top entry taking precedence. The reason for this terminology is that if the top line consists of the the integers $1, \ldots, n$ in order, and the bottom line consists of a permutation of these letters, then this is one way of writing a permutation, the permutation mapping the numbers of the top line to their partners on the bottom. The set of all generalized permutations is denoted by GP.

If π is a generalized permutation, we will let $\mathrm{top}(\pi)$ denote the row consisting of the entries in the top row of π. It is a vector whose entries are positive integers which are non-decreasing from left to right. So we can think of it as a one rowed tableau. If two adjacent entries of $\mathrm{top}(\pi)$ are equal, then the corresponding entries of $\mathrm{bot}(\pi)$ are non-decreasing from left to right. But otherwise there is no restriction on the bottom entries.

6.2 The bumping algorithm for generalized permutations.

Consider the word formed by $\mathrm{bot}(\pi)$ and construct the tableau corresponding to this word by the Shensted bumping algorithm, but every time a new box is created, put the corresponding entry of $\mathrm{top}(\pi)$ into this new box. At each stage we produce a pair of tableaux of the same shape. To illustrate, suppose that

$$\pi = \begin{array}{ccccccc} 1 & 1 & 1 & 2 & 2 & 3 & 7 \\ 2 & 3 & 3 & 1 & 2 & 1 & 1 \end{array}.$$

The word of bot(π) is 2331211. After the first three insertions we get the diagram

$$2\ 3\ 3$$

and will have place 1 in each of these three positions and so get the pair of diagrams

$$2\ 3\ 3 \qquad 1\ 1\ 1.$$

At the next stage we will insert 1 into the row 2 3 3 to obtain the diagram

$$\begin{array}{ccc} 1 & 3 & 3 \\ 2 & & \end{array}$$

and also place 2 in the new box of the second diagram so as to obtain the pair of diagrams

$$\begin{array}{ccc} 1 & 3 & 3 \\ 2 & & \end{array} \qquad \begin{array}{ccc} 1 & 1 & 1 \\ 2 & & \end{array}.$$

At the next stage we insert 2 into the first diagram and place 2 in the new box to obtain the pair

$$\begin{array}{ccc} 1 & 2 & 3 \\ 2 & 3 & \end{array} \qquad \begin{array}{ccc} 1 & 1 & 1 \\ 2 & 2 & \end{array}.$$

The next stage yields

$$\begin{array}{ccc} 1 & 1 & 3 \\ 2 & 2 & \\ 3 & & \end{array} \qquad \begin{array}{ccc} 1 & 1 & 1 \\ 2 & 2 & \\ 3 & & \end{array}.$$

The final stage yields

$$\begin{array}{ccc} 1 & 1 & 1 \\ 2 & 2 & 3 \\ 3 & & \end{array} \qquad \begin{array}{ccc} 1 & 1 & 1 \\ 2 & 2 & 7 \\ 3 & & \end{array}.$$

The bumping rule ensures that the left hand result is a tableau. Also the right hand result is a tableau because every time we add a box on the same row the number we are placing in the new box lies to the right of all preceding entries of top(π) which is non-decreasing from left to right. Also no two equal entries of top(π) can end up in the same column, because the lexicographic order insures that the corresponding entries on the bottom row are non-decreasing, and then our result about bumping routes says that the bumping route for the second entry lies strictly to the right of the route for the first. So both are tableaux, and we have shown how to associate to each generalized permutation a pair of tableaux of the same shape.

The second tableau can be used to keep track of how we performed the insertions. We can think of it a recording tableau for the insertions. For example, the location of the 7 in the right hand tableau above tells us that this box was where the last box created, and this allows us to reconstruct the bumping route and determine what the preceding left hand tableau was, and what was inserted. The next to the last box created must have been at the location of the rightmost 3 and so we can reconstruct its bumping route and so on. In short, the procedure is reversible.

6.3 The Robinson-Shensted-Knuth correspondence and Cauchy's formula.

We have constructed a bijective map between generalized permutations and pairs of tableau of the same shape. For any generalized permutation π we can calculate $\omega(\mathrm{top}(\pi))$ and $\omega(\mathrm{bot}(\pi)$. Introduce variables $X = (X_1, X_2, \dots)$ and $Y = (Y_1, Y_2, \dots)$. Associate a monomial in both variables to an generalized permutation by the rule

$$\pi \mapsto X^{\omega(\mathrm{top}(\pi))} Y^{\omega(\mathrm{bot}(\pi))} := \mathrm{mon}(\pi).$$

If the column $\begin{pmatrix} i \\ j \end{pmatrix}$ occurs k times in π, then it gives a contribution $X_i^k Y_j^k$ to the above monomial. So the generating function for mon is given by

$$\sum_{\pi \in GP} \mathrm{mon}(\pi) = \prod_{i,j \geq 1} \sum_{k \geq 0} X_i^k Y_j^k = \prod_{i,j \geq 1} \frac{1}{1 - X_i Y_j}$$

our old friend. On the other hand, for a pair (t, u) of tableau of the same shape, we can consider the monomial

$$X^{\omega(t)} Y^{\omega(u)}.$$

Summing over all tableaux of shape λ this gives $s_\lambda(X) s_\lambda(Y)$. The R-S-K correspondence preserves weights so we get Cauchy's formula (11).