Math. 126 Lecture 14

Symmetric Functions

November 18, 2002

Contents

1 Preliminaries and notation.

Let R be a commutative ring, and $X_1, X_2, \cdots, X_n$ a set of indeterminates. Let $\alpha = (\alpha_1, \ldots, \alpha_n)$ be a vector with non-negative integer coefficients. The set of all α form a monoid under addition which is isomorphic to the monoid of all monomials X^α under multiplication:

$$X^\alpha \cdot X^\beta = X^{\alpha+\beta}.$$

Here

$$X^\alpha = X_1^{\alpha_1} \cdots X_n^{\alpha_n}.$$

The corresponding algebra is the algebra of polynomials with coefficients in R usually denoted by $R[X_1, \ldots, X_n]$. We shall mainly be interested in the cases $R = \mathbb{Z}$ or $R = \mathbb{C}$ but will also have to consider $R = \mathbb{Q}$, the field of rational numbers.

We write

$$|\alpha| = \deg X^\alpha := \alpha_1 + \cdots + \alpha_n.$$

An element of $R[X_1, \ldots, X_n]$ of the form

$$f(X_1, \ldots, X_n) = \sum_{|\alpha|=d} c_\alpha X^\alpha$$

is called a homogenous polynomial of degree d, and the set of all such polynomials is denoted by $R[X_1, \ldots, X_n]^d$. This makes $R[X_1, \ldots, X_n]$ into a graded ring:

$$R[X_1, \ldots, X_n] = \bigoplus_{d=0}^{\infty} R[X_1, \ldots, X_n]^d,$$

$$R[X_1, \ldots, X_n]^r \cdot R[X_1, \ldots, X_n]^s = R[X_1, \ldots, X_n]^{r+s}.$$

We put a partial order on the set of all α (and hence on the set of all monomials) by declaring that $\alpha \prec \beta$ (so $X^\alpha \prec X^\beta$) if and only if

$$\sum_{i \leq k} \alpha_i \leq \sum_{i \leq k} \beta_i \quad \text{for all} \quad k \quad \text{and} \quad |\alpha| = |\beta|.$$

So monomials of different degrees are incomparable.

We may think of α as a function from $\{1, \ldots, n\}$ to $\mathbb{Z}^n$. Since the permutation group S_n acts on $\{1, \ldots, n\}$, it acts on the set of all α (and hence on the set of X^α):

$$w\alpha(i) = \alpha(w^{-1}(i)), \qquad w \in S_n.$$

Intuitively, w sends the variable X_i into the variable $X_{w(i)}$.

The orbit under S_n of any monomial X^α contains a unique element X^λ which is maximal for the partial order $\prec$, it is the monomial X^λ with $\lambda_1 \geq \cdots, \geq \lambda_n$. Such a λ is called a **partition** of $d = |\lambda|$ into n-parts (also known as a Young diagram with d boxes and $\leq n$ rows). The set of all partitions of d into n parts is denoted by $\mathcal{P}_{d,n}$. A partition of d (without qualification as to the number parts) is a non-increasing sequence of integers $\lambda_1 \geq \cdots \geq \lambda_m \geq \cdots$ such that $\sum \lambda_i = d$, so in particular $\lambda_i = 0$ for $i > d$. We denote the set of all partitions of d by $\mathcal{P}_d$ and regard $\mathcal{P}_{d,n}$ as a subset of $\mathcal{P}_d$. We set

$$\mathcal{P} = \bigcup_d \mathcal{P}_d.$$

It is the set of all Young diagrams. More geometrically, we can consider the Young diagram $Y(\lambda)$ to consist of all positions (i, j) (using matrix notation) where $j = 1, \ldots, \lambda_i$. Until the last two sections of this lecture, n will be fixed and all diagrams will be assumed without further notification to have at most n-rows.

The action of S_n of $R[X_1, \ldots, X_n]$ preserves degrees. A polynomial, i.e. an element of $R[X_1, \ldots, X_n]$, is called **symmetric** if it invariant under the action of S_n. We denote the set of symmetric polynomials of degree d in n-variables by Λ_n^d and set

$$\Lambda_n = \bigoplus_d \Lambda_n^d.$$

It is a subring of $R[X_1, \ldots, X_n]$ known as the ring (or algebra) of symmetric polynomials.

The simplest symmetric polynomials to write down are the **symmetric monomials** (bad but standard terminology -"symmetrized monomial" would be better) defined by

$$m_\lambda = m_\lambda(n) := \sum_{\alpha \in \text{the } S_n \text{ orbit of } \lambda} X^\alpha.$$

For example,

$$\begin{array}{rcl}
m_1 & = & X_1 + \cdots + X_n \\
m_2 & = & X_1^2 + \cdots + X_n^2 \\
m_{1,1} & = & X_1X_2 + X_1X_3 + \cdots + X_1X_n + X_2X_3 + \cdots = \displaystyle\sum_{1 \le i < j \le n} X_iX_j \\
m_3 & = & \displaystyle\sum_i X_i^3 \\
m_{2,1} & = & \displaystyle\sum_{i \ne j} X_i^2 X_j \\
m_{1,1,1,} & = & \displaystyle\sum_{i<j<k} X_iX_jX_k \\
\vdots & \vdots & \vdots
\end{array}$$

The m_λ are clearly independent. Also, if a homogenous polynomial of degree d, $f = \sum_{|\alpha|=d} c_\alpha X^\alpha$ is to be symmetric, then the c_α must be constant as α ranges over an S_n orbit. Thus the $m_\lambda, |\lambda| = d$ span Λ_n^d. So they form a basis of Λ_n^d. We will be interested in a different basis, the Schur basis, which we will now study.

2 Schur polynomials, skew Schur polynomials, and the Bender Knuth involution.

2.1 Skew Schur polynomials.

Let μ and ν be partitions with at most n parts, and suppose that $\nu_i \leq \mu_i$ for all i. This means that the diagram of ν is contained in the diagram of ν:

$$Y(\nu) \subset Y(\mu).$$

Let

$$D(\mu, \nu) = Y(\mu) \setminus Y(\nu).$$

So $D(\mu, \nu)$ consists of all pairs (i, j) with $1 \leq i \leq n$ and $\nu_i < j \leq \mu_i$. Let $\mathcal{S}(\mu/\nu)$ be the set of all (semi-standard) tableaux of shape $D(\mu, \nu)$ that is all maps $T : D(\mu, \nu) \to \{1, \ldots, n\}$ which are non-decreasing along rows (from left to right) and strictly decreasing down columns. Define the **weight** of T, denoted by $\omega(T) = (\omega_1(T), \ldots, \omega_n(T))$ where $\omega_k(T) = \#(T^{-1}k)$. So $\omega_k(T)$ is the number of times that k occurs in the diagram T. The **skew Schur polynomial** $s_{\mu/\nu}$ is defined by

$$s_{\mu/nu} := \sum_{T \in \mathcal{S}(\mu,\nu)} x^{\omega(T)}.$$

Our first task is to show that this is a symmetric polynomial.

Let s_k denote the involution $(k, k+1)$. (The letter s is getting overworked but I am following standard notation .) So $s_k \in S_n$ and

$$s_k(\alpha_1, \ldots, \alpha_{k-1}, \alpha_k, \alpha_{k+1}, \alpha_{k+2}, \ldots, \alpha_n) = (\alpha_1, \ldots, \alpha_{k-1}, \alpha_{k+1}, \alpha_k, \alpha_{k+2}, \ldots, \alpha_n).$$

To prove that $s_{\mu/\nu}$ is a symmetric polynomial, it is enough to show that there is an involution σ_k (i.e. a one to one map whose square is the identity) on $\mathcal{S}(\mu/\nu)$ such that

$$\omega(\sigma_k(T)) = s_k(\omega(T)).$$

2.2 The Bender Knuth involution.

Define an entry k in T to be *married* if there is a $k+1$ in the same column. Similarly define an entry $k+1$ to be married if there is a k in the same column. If k is married, its spouse $k+1$ must occur in the row directly beneath it since entries in a column are strictly increasing going down any column. If two entries k on the same row are part of a married couple, then every entry between them must be a married k, since entries are non-decreasing on each row. No entry k or $k+1$ on a row which contains a married k can occur to the left of all the married k's: certainly a $k+1$ can not occur there since the rows are non-decreasing. But also a k can not occur there: there must be a box below this position since there are boxes on the row beneath and to the right of this position, and the entry must be $> k+1$ since the k is unmarried, and this is impossible since the entries on the row beneath must be non-decreasing. Similarly, no unmarried k

or $k+1$ can occur to the right of a married $k+1$. So the set of all unmarried k's and $(k+1)$'s in any row form a contiguous sequence of positions. Above an unmarried k or $k+1$ if there is a position, it is filled by a number $< k$ and in the row below, if there is a position, it is filled by a number $> k+1$. Suppose that there are r unmarried k's followed by s unmarried $(k+1)$'s on a given row (where r or s or both might be 0). Replace this by a sequence of s k's followed by r $(k+1)$'s, and do this for each row. This is the desired involution σ_k invented by Bender and Knuth. Since the number of married k's is the same as the number of married $(k+1)$'s, we have the desired equality $\omega(\sigma_k(T)) = s_k(\omega(T))$.

2.3 The Schur polynomials.

When $\nu = 0$ we write s_μ instead of $s_{\mu/0}$. So

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

where the sum is over all tableau of shape λ. We now know that this is a symmetric polynomial.

In designing a tableau of shape λ the largest number of ones that we can insert is obtained by filling the whole top row with ones; so $\omega_1(T) \leq \lambda_1$ and λ_1 is achieved by filling the top row with ones. Then the maximum value of $\omega_1(T) + \omega_2(T)$ is achieved by filling the entire second row with 2's etc. In other words, if T is of shape λ then

$$\omega(T) \prec \lambda$$

with equality for the unique tableau obtained by filling the i-th row with i. (Let us call this tableau the canonical tableau of shape λ.) The argument shows (since s_λ is symmetric) that

$$s_\lambda = m_\lambda + \cdots$$

where the remaining terms involve monomials which are $\prec \lambda$, and symmetric, and hence are linear combinations of μ with $\mu \prec \lambda,\ \mu \neq \lambda$. In other words, if we write out the expansion of the s_λ in terms of the m_ν, we get an "upper triangular matrix" relative to the ordering $\prec$ with ones on the diagonal (and non-negative integers above the diagonal). This matrix is invertible over the integers, and this shows that the s_λ form a basis of Λ_n^d as the λ range over all diagrams with d boxes and at most n rows.

We will return to make use of the fact that the canonical tableau is the unique tableau with shape λ and whose weight is a partition.

3 Stembridge's theorem.

For any $\beta = (\beta_1, \ldots, \beta_n) \in \mathbb{Z}^n, \beta_i \geq 0$, we have defined X^β as $X^\beta = X_1^{\beta_1} \cdots X_n^{\beta_n}$. Define the **alternant** a_β as

$$a_\beta := \det(X_i^{\beta_j}) = \sum_{u \in S_n} \operatorname{sgn}(u) X^{u\beta}.$$

For any $w \in S_n$ we have

$$a_{w\beta} = \operatorname{sgn} w a_\beta.$$

In what follows let μ and ν be partitions with $\nu_i \leq \mu_i$ for all i, so that we have μ/ν, $\mathcal{S}(\mu/\nu)$, and $s_{\mu/\nu}$, as above. Also set

$$\rho := (n-1, n-2, \ldots, 0)$$

and

$$\mathbf{0} = (0, \ldots, 0).$$

For any $T \in \mathcal{S}(\mu/\nu)$ we let $T_{\geq j}$ denote the tableau obtained formed by entries in the columns $j, j+1, \ldots$ of T. In other words it is the portion of T to the right of and including the j-th column. We use similar notations such as $T_{<j}, T_j$ and $T_{>j}$. Let λ be an element of $\mathcal{P}_n$. We say that T is a Good Guy (relative to λ) if for all j, $\lambda + \omega(T_{\geq j})$ is a partition, i.e. its entries are decreasing. The other $T \in \mathcal{S}(\mu/\nu)$ are Bad Guys.

3.1 Stembridge's theorem.

Stembridge's theorem asserts that

$$a_{\lambda+\rho} s_{\mu/\nu} = \sum a_{\lambda+\omega(T)+\rho}$$

where the sum is over Good Guys.

Proof. We know (from the Bender Knuth involution) that for any $w \in S_n$ there are just as many T of weight ω as there are of weight $w\omega$ as T ranges over $\mathcal{S}(\mu/\nu)$. Hence the quantities $w(\lambda + \rho) + \omega(T)$ and $w(\lambda + \rho + \omega(T))$ occur the same number of times as T ranges over $\mathcal{S}(\mu/\nu)$. Therefore

$$a_{\lambda+\rho} s_{\mu/\nu} = \sum_{w \in S_n} \sum_{T \in \mathcal{S}(\mu/\nu)} \operatorname{sgn}(w) X^{w(\lambda+\rho)} X^{\omega(T)} =$$

$$\sum_{w \in S_n} \sum_{T \in \mathcal{S}(\mu/\nu)} \operatorname{sgn}(w) X^{w(\lambda+\rho+\omega(T))} = \sum_{T \in \mathcal{S}(\mu/\nu)} a_{\lambda+\omega(T)+\rho}.$$

Here the sum is over all $T \in \mathcal{S}(\mu/\nu)$. To prove the theorem we must show that the contributions of the Bad Guys cancel out. To see the idea, suppose that $\lambda_k = \lambda_{k+1}$ and that T has no k's and exactly one $k+1$. Then T is a Bad Guy because $\lambda+\omega(T)$ has $\lambda_{k+1}+\omega_{k+1}(T)=\lambda_k+\omega_k(T)+1$ which is illegal. But when we add the ρ we have $\rho_{k+1}=\rho_k-1$ so $\lambda_{k+1}+\omega_{k+1}(T)+\rho_{k+1}=\lambda_k+\omega_k(T)+\rho_k$ and so the corresponding $a_{\lambda+\omega(T)+\rho}=0$ since two successive columns in the determinant defining $a_{\lambda+\omega(T)+\rho}$ are equal.

Now let us turn to the general proof. Suppose T is a Bad Guy, which means that $\lambda+\omega(T_{\geq j})$ fails to be a partition for some j. Choose j to be the largest for which this holds, and among those k, j for which this holds choose k to be the smallest. So

$$\lambda_k+\omega_k(T_{\geq j})-(\lambda_{k+1}+\omega_{k+1}(T_{\geq j}))<0$$

while

$$\lambda_k+\omega_k(T_{>j})-(\lambda_{k+1}+\omega_{k+1}(T_{>j}))\geq 0.$$

Each column of T contains at most one k and at most one $k+1$. So the only way that this can happen is if the last inequality is an equality and the j-th column of T contains $k+1$ but no k, and then

$$\lambda_{k+1}+\omega_{k+1}(T)=\lambda_k+\omega_k(T)+1.$$

The entry above the $k+1$ in the j-th column of T is strictly less than k, and so all the entries of the row above the $k+1$ in the j-th column and to the left of that column are also strictly less than k. So if we apply the Bender Knuth involution to $T_{<j}$ leaving $T_{\geq j}$ unchanged, we still have a tableau. So let T^* denote this tableau obtained from T by applying the Bender Knuth involution to $T_{<j}$ leaving $T_{\geq j}$ unchanged. Since we have not changed $T_{\geq j}$ we have that T^* is still a Bad Guy, and the map $T\mapsto T^*$ is an involution on the set of Bad Guys. Also, $s_k\omega(T_{<j})=\omega(T^*_{<j})$ while $\lambda_{k+1}+\omega_{k+1}(T)=\lambda_k+\omega_k(T)+1$. This implies that

$$s_k(\lambda+\omega(T)+\rho)=\lambda+\omega(T^*)+\rho$$

and hence that

$$a_{\lambda+\omega(T)+\rho}=-a_{\lambda+\omega(T^*)+\rho}$$

showing that the contributions of the Bad Guys cancel out. QED

3.2 The bi-alternant formula aka as the Weyl character formula.

Suppose we take $\lambda=\nu=\mathbf{0}$. Then there is only one tableau of shape μ which is a Good Guy, i.e. which is a partition, and it is the canonical tableau of shape μ obtained by putting all ones in the top row of μ etc. It has weight μ. So Stembridge's theorem yields

$$a_\rho s_\mu=a_{\mu+\rho}$$

or

$$s_\mu=\frac{a_{\mu+\rho}}{a_\rho}.$$

This the bi-alternant formula. It is a special case of the Weyl character formula for representations of compact Lie groups.

3.3 Zelevinsky's rule.

If we substitute

$$s_\lambda = \frac{a_{\lambda+\rho}}{a_\rho}$$

into Stembridge's theorem we get

$$s_\lambda s_{\mu/\nu} = \sum s_{\lambda+\omega(T)}$$

where the sum is over Good Guys. This is Zelevinsky's rule.

3.4 Expansion of skew Schur functions.

We know that $s_{\mu/\nu}$ is a symmetric polynomial, and we know that the ordinary Schur functions form a basis of the symmetric polynomials. We can ask for the coefficients of the expansion of $s_{\mu/\nu}$ in terms of this basis. Taking $\lambda = \mathbf{0}$ in Zelevinsky's rule gives us this expansion.

3.5 The Littlewood-Richardson rule.

Taking $\nu = \mathbf{0}$ in Zelevinky's rule gives us a formula for the product of two Schur functions in terms of the Schur function basis. This is the Littlewood-Richardson rule. The Pieri formulas are a special case.

Here is a link (kindly sent to me by Prof. Stembridge) to a Maple program which implements the LR rule:

http://www.math.lsa.umich.edu/~jrs/software/SFexamples/LR_rule

He also suggests that you visit the web site

http://www.math.lsa.umich.edu/~jrs/maple.html

.

4 Three other bases.

There are three other famous bases of Λ_n^d:

4.1 The elementary symmetric functions.

The Schur function corresponding to a column with k rows is called the k-th elementary symmetric function, and is denoted by e_k. So

$$\begin{array}{rcl} e_1 & = & X_1 + \cdots + X_n \\ e_2 & = & \displaystyle\sum_{1 \le i < j \le n} X_i X_j \\ \vdots & \vdots & \vdots \\ e_k & = & \displaystyle\sum_{1 \le i_1 < \cdots < i_j \le n} X_{i_1} \cdots X_{i_j} \\ \vdots & \vdots & \vdots \\ e_n & = & X_1 \cdots X_n. \end{array}$$

Then for any diagram λ define

$$e_\lambda := e_{\lambda_1} \cdots e_{\lambda_n}.$$

Then the e_λ form a basis of Λ_n^d as the λ range over diagrams with d boxes and at most n *columns*. To see this, consider for example $e_{4,1} = e_4 \cdot e_1$. The leading monomial in this product is the leading monomial in

$$(X_1 X_2 X_3 X_4 + \cdots)(X_1 + \cdots)$$

which is $X_1^2 X_2 X_3 X_4$ which is X^μ where $\mu = (2, 1, 1, 1)$. The diagram μ is the "dual diagram" to λ obtained by flipping λ over the diagonal so that the columns of λ become the rows of μ. This is true in general - the leading term in e_λ is x^μ where μ is dual to λ as can readily be checked. So we have an "upper triangular matrix" with ones on the diagonal expressing the e_λ in terms of the m_μ. This proves that the elementary symmetric functions e_λ of degree d (with at most n-columns) form a basis of Λ_n^d. This also shows that $\Lambda_n = R[e_1, \ldots, e_n]$ as an algebra - an assertion which is known as the fundamental theorem of symmetric polynomials.

We should also note that

$$\prod_{i=1}^{n} (1 - tX_i) = \sum_{q=1}^{n} (-1)^q e_q t^q,$$

We say that $\prod_{i=1}^{n}(1 - tX_i)$ is the generating function of $(-1)^q e_q$. I call this the Fermi-Dirac function.

4.2 The complete symmetric functions.

This is the same as the above except that you start with the Schur functions of a row with at most n columns. They will be the key for translating the

Littlewood-Richardson rule that we have proved for symmetric functions into a rule for induced representations of the symmetric group.

We have

$$\prod_{i=1}^{n} \frac{1}{1 - X_i t} = \sum h_p t^p$$

as the generating function for complete elementary symmetric functions $h_p = s_{(p)}$. I call this the Bose-Einstein picture. Notice that it follows from

$$\prod_{i=1}^{n}(1 - tX_i) \cdot \prod_{i=1}^{n} \frac{1}{1 - X_i t} = 1$$

that

$$h_k - e_1 h_{k-1} + e_2 h_{k-2} - \cdots + (-1)^k e_k = 0.$$

This means that we can express the h's in terms of the e's and vice versa, so the algebra $R[h_1, \ldots, h_n]$ is the same as the algebra $R[e_1, \ldots, e_n]$ which is the algebra of symmetric polynomials. So the h_λ as λ ranges over all diagrams with d boxes and at most n columns forms a basis for Λ_n^d.

4.3 Power sums.

The four bases we have studied so far work over any ring since they work over the integers. We will now introduce a fifth collection of bases for Λ_n^d which will require division by an integer, and so only work if our ring contains the rational numbers $\mathbb{Q}$. For any non-negative integer r define

$$p_r := X_1^r + \cdots + X_n^r.$$

These are known as the **Newton power sums**. These are clearly symmetric and then for any diagram λ with k rows we define

$$p_\lambda := p_{\lambda_1} \cdot p_{\lambda_2} \cdots p_{\lambda_k}.$$

If we take the logarithmic derivative of

$$E(t) := \prod_{i=1}^{n}(1 + X_i t) = \sum e_p t^p$$

we obtain

$$\frac{d}{dt}(\log E(t)) = \frac{E'(t)}{E(t)} = \sum_i \frac{d}{dt}\log(1 + X_i t) = \sum_i \frac{1}{1 + X_i t} = \sum_{q=0}^{\infty}(-1)^q p_q t^q.$$

So

$$(e_1 + 2e_2 t + 3e_3 t^2 + \cdots) = (e_0 + e_1 t + e_2 t^2 + \cdots)(1 - p_1 + p_2 t - p_3 t^2 + \cdots).$$

Comparing coefficients shows that for every k we have

$$ke_k - p_1 e_{k-1} + p_2 e_{k-2} + \cdots + (-1)^k p_k = 0.$$

This proves that if our ring contains the rational numbers then

$$R[p_1, \ldots, p_k] = R[e_1, \ldots, e_k].$$

5 Symmetric functions.

In the lectures that follow it is going to be a pain to specify the number of variables. One way of dealing with this is to wimp out and always choose n "sufficiently large". Another way is to introduce the ring of formal power series in infinitely many variables. I will follow the compromise path chosen by Fulton; Notice that all the kinds of symmetric polynomials we have been considering *specialize* in the sense that if P is a polynomial in n variables then

$$P(x_1, \ldots, x_\ell, 0, \ldots, 0) = P(x_1, \ldots, x_\ell),$$

where, on the right hand side we are considering the corresponding polynomial in fewer variables. So let us define a **symmetric function of degree** d to be a collection of homogenous polynomials of degree d in an increasing number of variables which satisfy the above consistency condition. We let Λ^d be the module (over the ring R) of all such collections with coefficients in R. So for each partition λ of d the elements $m_\lambda, s_\lambda, e_\lambda, h_\lambda$ and p_λ form bases of Λ^d. Also multiplication extends consistently so

$$\Lambda := \bigoplus_d \Lambda^d$$

is a ring over R. It can be identified with the ring of polynomials in the infinitely many variables $e_1, e_2, \ldots$ or $h_1, h_2, \ldots$ and the Littlewood-Richardson rule for the multiplication of Schur functions holds in this algebra.

Our next task will be to translate facts about the algebra of symmetric functions (in particular the Littlewood-Richardson rule) into statements about representations of symmetric groups.