


NOTES FOR SEPTEMBER 7

KEVIN CARDE

1. TERMINOLOGY FROM LAST TIME

Recall from last time the ring of symmetric polynomials $\Lambda_n = \mathbb{Z}[x_1, \dots, x_n]^{S_n}$ and the ring of symmetric Laurent polynomials $\Lambda_n^\pm = \mathbb{Z}[x_1^\pm, \dots, x_n^\pm]^{S_n}$.

A **partition** is a finite sequence $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$ of weakly decreasing nonnegative integers (with trailing zeroes usually ignored). We draw partitions as **Young diagrams** such that the partition $(4, 2, 1)$ is represented as



(this is the English way to draw partitions—there are other conventions).

The **size** of a partition is the sum of its parts: $|\lambda| = \lambda_1 + \lambda_2 + \dots + \lambda_n$, so $|(4, 2, 1)| = 7$. The **length** of a partition is the number of nonzero rows, so $\ell(4, 2, 1) = 3$.

There is a basis of Λ_n indexed by partitions λ with $\ell(\lambda) \leq n$ called the **monomial basis**. For example,

$$m_{111}(x_1, x_2, x_3, x_4) = x_1x_2x_3 + x_1x_2x_4 + x_1x_3x_4 + x_2x_3x_4.$$

In general, m_λ is the sum of all ways to use the parts of the partitions as exponents on the variables. The degree of m_λ is $|\lambda|$, and the dimension of the degree d part of Λ_n is the number of partitions of size d with at most n rows.

Recall also **majorization order**: $\lambda \preceq \mu$ if

$$\begin{aligned} \lambda_1 &\leq \mu_1 \\ \lambda_1 + \lambda_2 &\leq \mu_1 + \mu_2 \\ &\vdots \\ \lambda_1 + \lambda_2 + \dots + \lambda_n &\leq \mu_1 + \mu_2 + \dots + \mu_n. \end{aligned}$$

2. ELEMENTARY SYMMETRIC POLYNOMIALS

Define $e_k = m_{11\dots 1}$, where the subscript on m consists of k 1s. For example, $e_0 = m_0 = 1$ and $e_3(w, x, y, z) = wxy + wxz + wyz + xyz$. Then set

$$e_\lambda = e_{\lambda_1} e_{\lambda_2} \dots e_{\lambda_n}$$

so that e.g.

$$\begin{aligned} e_{31}(w, x, y, z) &= e_3(w, x, y, z)e_1(w, x, y, z) \\ &= (wxy + wxz + wyz + xyz)(w + x + y + z) \\ &= (w^2xy + \dots) + 4wxyz \\ &= m_{211} + 4m_{1111}. \end{aligned}$$

These are called the **elementary symmetric polynomials**.

Lemma 1. $e_\lambda = m_{\lambda^T} + (\text{linear combination of } m_\mu \text{ for } \mu \preceq \lambda^T)$.

(Short rant: If we denoted by “ e_λ ” what’s standardly called “ e_{λ^T} ”, there would be a lot fewer transposes in the theory! Unfortunately, this way is totally standard.)

Proof. Notice that the largest exponent of x_1 in e_λ is at most $\ell(\lambda) = (\lambda^T)_1$. Notice that the largest sum of the exponents on x_1 and x_2 is at most the number of parts of λ plus the number of parts of λ bigger than 1, or $(\lambda^T)_1 + (\lambda^T)_2$. In general,

$$\sum_{k=1}^r (\text{exponent of } x_k) \leq (\lambda^T)_1 + (\lambda^T)_2 + \cdots + (\lambda^T)_r.$$

Therefore, if m_μ shows up when expanding e_λ in the monomial basis, then $\mu \preceq \lambda^T$. \square

Corollary 2. *The e_λ with $\ell(\lambda^T) \leq n$ are a basis for Λ_n .*

Proof. Fix a degree d ; we will show that the appropriate e_λ form a basis of the degree d part of Λ_n . Let $X = \{e_\lambda : \ell(\lambda^T) \leq n, |\lambda| = d\}$ and let $Y = \{m_\lambda : \ell(\lambda) \leq n, |\lambda| = d\}$. We have seen that Y is a basis of the degree d part of Λ_n ; our goal now is to show that X is also such a basis.

Clearly $|X| = |Y|$, so there are the correct number of e_λ to be a basis. Therefore, since the degree d part of Λ_n is finite dimensional, it suffices to show that we can express all the $m_\mu \in Y$ in terms of the $e_\lambda \in X$.

Extend the majorization order to a total order (for example, take the lexicographic order). Order the elementary symmetric polynomials in X so that e_λ comes before e_μ whenever $\lambda^T \leq \mu^T$ in this total order, and order the monomial symmetric polynomials in Y so that m_λ comes before m_μ whenever $\lambda \leq \mu$. Let E be the (column) vector of elementary symmetric polynomials from X in this order from largest to smallest, and let M be the (column) vector of monomial symmetric polynomials from largest to smallest.

Let $a_{\lambda\mu}$ be the coefficient of m_μ in e_{λ^T} . Let $A = (a_{\lambda\mu})$ be the matrix of these $a_{\lambda\mu}$ (with rows and columns from largest to smallest in our total order on partitions). Hence the matrix A takes the vector M to the vector E : $AM = E$.

Notice now that this matrix is upper triangular, as the previous lemma tells us that $a_{\lambda\mu}$ is 0 unless $\mu \preceq \lambda$, and further, this matrix has 1s on the diagonal (since $a_{\lambda\lambda} = 1$). Hence this matrix is invertible (and the inverse is once again an upper triangular matrix with 1s on the diagonal). Therefore, we have $M = A^{-1}E$, which, when expanded out, shows that we can successfully write all the $m_\mu \in Y$ in terms of the $e_\lambda \in X$, as desired. \square

Example 3. This kind of argument is important, but in this case it's unfortunately complicated by the pesky transpose issues. To (hopefully) make this clearer, let's work through the proof for $n = 3$ and $d = 4$.

There are 5 partitions of 4, all but one of which have at most 3 parts. Hence our set Y consists of $\{m_4, m_{31}, m_{22}, m_{211}\}$, which we already know to be a basis of the degree 4 part of Λ_3 . Our set X consists of all partitions of 4 whose *transpose* has at most 3 parts; equivalently, all partitions of 4 with largest part at most 3. So we have $X = \{e_{1111}, e_{211}, e_{22}, e_{31}\}$.

The majorization order on partitions of 4 is already a total order, with $(4) \succeq (3, 1) \succeq (2, 2) \succeq (2, 1, 1) \succeq (1, 1, 1, 1)$, and our sets Y and X as written above are already ordered appropriately (recall that X is ordered according to the majorization order of the transposes, which is equivalent (homework!) to reverse majorization). In accordance with the lemma, we have

$$\begin{aligned} e_{1111} &= (x + y + z)^4 = m_4 + 4m_{31} + 6m_{22} + 12m_{211} \\ e_{211} &= (yz + xz + xy)(x + y + z)^2 = m_{31} + 2m_{22} + 5m_{211} \\ e_{22} &= (yz + xz + xy)^2 = m_{22} + 2m_{211} \\ e_{31} &= xyz(x + y + z) = m_{211} \end{aligned}$$

and so our matrix equation is

$$\begin{pmatrix} 1 & 4 & 6 & 12 \\ 0 & 1 & 2 & 5 \\ 0 & 0 & 1 & 2 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} m_4 \\ m_{31} \\ m_{22} \\ m_{211} \end{pmatrix} = \begin{pmatrix} e_{1111} \\ e_{211} \\ e_{22} \\ e_{31} \end{pmatrix}.$$

This matrix is upper triangular with ones on the diagonal, as promised, and inverting it allows us to write the monomial symmetric polynomials in terms of the elementary symmetric polynomials.

As a corollary, we get what is sometimes called the Fundamental Theorem of Symmetric Polynomials:

Corollary 4. $\Lambda_n = \mathbb{Z}[e_1, e_2, \dots, e_n]$.

Proof. Monomials in e_1, \dots, e_n are exactly the e_λ with all parts of λ less than n , or the e_λ with $\ell(\lambda^T) \leq n$. \square

If $\ell(\lambda) \leq n$, then the expression for m_λ as a polynomial in the e_k does not depend on n . We have ring maps

$$\begin{aligned} \Lambda_n &\rightarrow \Lambda_{n-1} \\ x_n &\mapsto 0 \\ x_k &\mapsto x_k, \text{ for } k < n \end{aligned}$$

or in terms of the e_k ,

$$\begin{aligned} e_n &\mapsto 0 \\ e_k &\mapsto e_k, \text{ for } k < n. \end{aligned}$$

Therefore, if we set $\Lambda = \mathbb{Z}[e_1, e_2, \dots, e_n, \dots]$, we can talk about $m_\lambda \in \Lambda$ as the polynomial in the e_k which works in Λ_n for all n (taking $e_k = 0$ whenever $k > n$). For example, we can talk about $x_1^2 + x_2^2 + x_3^2 + \dots$ in any number of variables as $e_1^2 - 2e_2$.

Another way to define this is to notice that the coefficients for multiplying the m_λ don't depend on n :

$$m_\lambda m_\mu = \sum a'_{\lambda\mu} m_\nu$$

where the $a'_{\lambda\mu}$ don't depend on n as soon as n is large enough to have m_ν in the first place, and the number of terms on the right hand side is at most the number of partitions of size $|\lambda| + |\mu|$. Hence we can say Λ has basis m_λ (for all partitions λ) with multiplication constants $a'_{\lambda\mu}$.

This approach allows us to define Λ^\pm analogously. Λ_n^\pm has basis m_λ with $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$, where now the λ_k may be less than zero. We see the same coefficients when multiplying together the m_λ regardless of n , and again we can define Λ^\pm to use these multiplication constants.

However, Λ^\pm does not have a presentation as nice as $\Lambda = \mathbb{Z}[e_1, e_2, \dots]$. For finite n , however, we can write $\Lambda_n^\pm = \mathbb{Z}[e_1, e_2, \dots, e_{n-1}, e_n^\pm]$. We will be working in Λ all the time because it's convenient, but it's worth being aware of these issues with Λ^\pm .