Math 821 Lecture Notes

Ross Churchley March 14 (Happy Pi Day!)

Jacobi triple product formula

Proposition 1.

$$\sum_{h=-\infty}^{\infty} x^{h^2} y^h = \prod_{i=1}^{\infty}(1-x^{2i-1}y)(1+x^{2i-1}y^{-1})(1-x^{2i})$$

Proof. It suffices to show

$$\left(\prod_{i=1}^{\infty}\frac{1}{1-x^{2i}}\right)\sum_{h=-\infty}^{\infty} x^{h^2} y^h = \prod_{i=1}^{\infty}(1-x^{2i-1}y)(1+x^{2i-1}y^{-1}). \tag{1}$$

The left-hand side consists of an infinite product counting partitions with even parts, multiplied by a sum generating integers as sides of squares. The product on the right-hand side counts partitions with odd and distinct parts in two ways: in the first bracket, y counts the numbers of parts, while y^{-1} takes the place of y in the second bracket.

By the next homework, partitions with odd and distinct parts are in bijection with self-conjugate partitions[1] in such a way that the number of parts in the partition with odd distinct parts maps to the number $d(\lambda)$ of boxes on the diagonal[2] in the self-conjugate partition. With this bijection, our goal becomes the following.

[1] a *self-conjugate partition* is one which is equal to its conjugate

[2] by the *diagonal* we mean the line $y=-x$, assuming the top left corner of the Ferrers diagram is at $(0,0)$

Let $\mathcal{P}_e$ be the combinatorial class of partitions with even parts and $\mathcal{S}$ be the combinatorial class of self-conjugate partitions. We will be done if we construct a bijection $\mathcal{P}_e \times \mathbb{Z} \to \mathcal{S} \times \mathcal{S}$ which preserves the powers of x and y in equation (1): i.e. if $(\mu, h) \mapsto (\lambda_1, \lambda_2)$,

$$|\lambda_1| + |\lambda_2| = |\mu| + h^2$$
$$d(\lambda_1) - d(\lambda_2) = h.$$

Take $(\mu, h) \in \mathcal{P}_e \times \mathbb{Z}$. First say $h \geq 0$. Since μ has all even parts, let ρ be the partition with each part half the corresponding part in μ. Now put an $h \times h$ square with the top left corner at $(0,0)$, with the Ferrers diagram of ρ with left corner at $(0,-h)$ and the Ferrers diagram of $\tilde{\rho}$ with the top left corner at $(h,0)$. In general, the result is not quite a Ferrers diagram (see Figure 1). However, we can derive two self-conjugate Ferrers diagrams from it as follows.

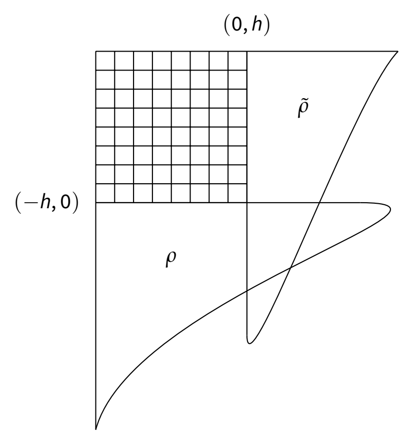

Figure 1: The not-quite-Ferrers-diagram used in the construction of α

Let α be the partition whose Ferrers diagram consists of the $h \times h$ square, the boxes of ρ's diagram which lie strictly below the diagonal $y=-x$, and the boxes of $\tilde{\rho}$ which lie on or above the diagonal. Let β by taking the boxes in this diagram that lie below the line $y=-h$ and to the right of $x=h$.

Our bijection $\mathcal{P}_e \times \mathbb{Z} \to \mathcal{S} \times \mathcal{S}$ will have $(\mu, h) \mapsto (\alpha, \beta)$ when $h \geq 0$. When $h < 0$, we perform the same construction with $|h|$ in place of h and have $(\mu, h) \mapsto (\beta, \alpha)$.

Let's check α and β are self-conjugate. Write $\alpha = (\alpha_1, \alpha_2, \ldots,)$ and $\beta = (\beta_1, \beta_2, \ldots,)$. From the diagram, we have

$$\alpha_i = \begin{cases} h + \tilde{\rho}_i & i \leq h \\ \min\{i, \rho_{i-h}\} + \max\{0, \tilde{\rho}_i - (i-h)\} & \text{otherwise} \end{cases}$$

but since conjugating the double diagram leaves it invariant, we have the same formula for $\tilde{\alpha}_i$. A similar argument shows that

$$\beta_i = \min\{\rho_i - h, \tilde{\rho}_{i+h}\} = \tilde{\beta}_i.$$

Count the number of boxes in the double-diagram in Figure 1, double-counting the boxes where ρ and $\tilde{\rho}$ overlap. On the one hand, this is $h^2 + 2|\rho| = h^2 + |\mu|$. On the other hand, since α and β overlap in the same boxes as ρ and $\tilde{\rho}$, this is $|\alpha| + |\beta|$. So

$$|\alpha| + |\beta| = |h|^2 + |\mu|.$$

We also have $d(\alpha) - d(\beta) = |h|$, as the boxes of β on the diagonal $y = -x$ are precisely the boxes of α on the diagonal to the right of $x = h$.

Now let's check this is a bijection by giving the inverse map. Take $(\lambda_1, \lambda_2) \in \mathcal{S} \times \mathcal{S}$. Take the Ferrers diagram of λ_1 and λ_2 and line them up so the square on the diagonal for each coincide. Let $h = d(\lambda_1) - d(\lambda_2)$. If $h \geq 0$, let $\alpha = \lambda_1$ and $\beta = \lambda_2$, otherwise let $\alpha = \lambda_2$ and $\beta = \lambda_1$. Let ρ be the partition with Ferrers diagram built of the parts of those squares of α below $y = -|h|$ and strictly below $y = -x$ along with those squares of β which are on or above $y = -x$, where the top corner of α is at $(0, 0)$. This construction is inverse to the original and hence we have a bijection. □

Symmetric functions

Let $\underline{x} = (x_1, x_2, \ldots)$ be a countable sequence of variables.

Definition. A *monomial* $\underline{x}^\alpha$ in the variables $\underline{x}$ indexed by $\alpha = (\alpha_1, \alpha_2, \ldots)$, $\alpha_i \in \mathbb{Z}_{\geq 0}$ is a product $\prod_{i=1}^{\infty} x_i^{\alpha_i}$ where only finitely many α_i are nonzero.

Definition. The *degree* of a monomial $\underline{x}^\alpha$ is $\sum_{i=1}^{\infty} \alpha_i$.

Definition. Let $R(\underline{x})$ be the ring of formal power series $\sum_\alpha c_\alpha \underline{x}^\alpha$ of bounded degree—i.e. for each element $\sum_\alpha c_\alpha \underline{x}^\alpha \in R(\underline{x})$, there is a d such that $deg(\underline{x}^\alpha) > d$ implies $c_\alpha = 0$.

Some examples:

$$\begin{aligned} &x_1 + x_2 + x_3 + \cdots && \in R(\underline{x}) \\ &x_1^2 x_2 + x_3^{85} && \in R(\underline{x}) \\ &x_1 + x_2^2 + x_3^3 + \cdots && \notin R(\underline{x}) \end{aligned}$$

Definition. The permutation group S_n acts on $R(\underline{x})$ by acting on the first n variables.

Definition. The ring of symmetric functions in $\underline{x}$ is $\Lambda(\underline{x}) = \{f \in R(\underline{x}) : f$ is invariant under the action of S_n for all $n\}$.

Definition. Under this action we can view $S_{n-1} \subseteq S_n$ and so let $S_\infty = \bigcup_{\geq 1} S_n$

Examples:

$x_1 + x_2 + x_3 + \cdots$	$\in \Lambda(\underline{x})$
the other above examples	$\notin \Lambda(\underline{x})$
$x_1^2 x_2 + x_2^2 x_1 + \cdots + x_{2013}^2 x_1 + \cdots$	$\in \Lambda(\underline{x})$

What can we say about the exponents appearing? The only information after permuting is the partition of the powers which appear.

Definition. Given a partition $\lambda = (\lambda_1, \ldots, \lambda_k)$ define $m_\lambda \in \Lambda(\underline{x})$ as follows: let $S(\lambda) = \{\sigma(\lambda_1, \lambda_2, \ldots, \lambda_k, 0, 0, \ldots, 0) : \sigma \in S_n, n \geq k\}$. (This is really just the S_∞ orbit of λ). Then $m_\lambda = \sum_{\alpha \in S(\lambda)} \underline{x}^\alpha$. These are called *monomial symmetric functions*.

Example 1. Let's work out m_λ for some partitions λ of 4:

$$m_{(4)} = x_1^4 + x_2^4 + x_3^4 + \cdots$$
$$m_{(3,1)} = x_1^3 x_2 + x_1 x_3^3 + \cdots$$

Proposition 2. *The m_λ form a vector space basis for $\Lambda(\underline{x})$.*

Proof. Given $f \in \Lambda(\underline{x})$, f is invariant under S_∞ so the powers appear in f consist of a disjoint union of some orbits under this action. Each orbit O corresponds to a partition. Finally, f is of bounded degree, so the partitions appearing are of bounded size, so there are finitely many of them. □

Do we really need infinitely many variables? No, but you need enough—three variables is enough for m_λ when λ is a partition of 3:

$$m_{(1,1,1)}|_{\Lambda(x_1,x_2,x_3)} = x_1 x_2 x_3,$$

but since

$$m_{(1,1,1)}|_{\Lambda(x_1,x_2)} = 0$$

two variables are not enough.

Formally, for all $n \geq 1$ there is an algebraic homomorphism $\phi_n : \Lambda(\underline{x}) \to \Lambda(x_1, x_2, \ldots, x_n$ mapping $x_i \mapsto x_i$ when $i \leq n$ and mapping all other x_i to 0. Futhermore, $\Lambda(\underline{x})$ is a graded vector space graded by degree and so is $\Lambda(x_1, \ldots, x_n)$.

Proposition 3. *$\phi_n : \Lambda(\underline{x}_i \to (\Lambda(x_1, \dots, x_n))_i$ is a vector space isomorphism for $i \leq n$.*

Proof. A basis for $\Lambda(\underline{x})_i$ is $\{m_\lambda : |\lambda| = i\}$ so we just need to check that their images under ϕ form a basis for $\Lambda(x_1, \dots, x_n)_i$. A partition λ of i has at most i parts, so there is at least one monomial in m_λ using at most the first i variables of $\underline{x}$. Since $i \leq n$ this gives $\phi(m_\lambda) \neq 0$. Further more, no monomial appears in more than one m_λ so there can be no cancellation. So the image of any linear combination of $\{m_\lambda : |\lambda| = i\}$ is also nonzero, so ϕ_n is one-to-one.

Since the m_λ span $\Lambda(\underline{x})_i$, so do the $\phi_n(m_\lambda)$. Hence they form a basis. □

Note that the proposition shows that any identity true in $\Lambda(x_1, \dots, x_n)$ for all n is also true in Λ: since any identity is finite, it must appear at some finite level of the grading.

References

Victor Reiner, *Hopf algebras in combinatorics* lecture notes (Chapter 2). `http://www.math.umn.edu/~reiner/Classes/HopfComb.pdf`