NOTES ON MACDONALD POLYNOMIALS AND THE GEOMETRY OF HILBERT SCHEMES

MARK HAIMAN (mhaiman@math.berkeley.edu) *
U.C. Berkeley

Abstract. These notes are based on a series of seven lectures given in the combinatorics seminar at U.C. San Diego in February and March, 2001. My lectures at the workshop which is the subject of this proceedings volume covered a portion of the same material in a more abbreviated form.

Key words: Symmetric functions, Macdonald polynomials, Hilbert schemes

2000 Mathematics Subject Classification: Primary 14C05; Secondary 05E05, 14F17, 14M05

1. History and introduction

1.1. OVERVIEW

In these lectures we'll be discussing a series of new results in combinatorics, algebra and geometry. The main combinatorial problems we solve are (1) we prove the *positivity conjecture* for Macdonald polynomials, and (2) we prove a series of conjectures relating the *diagonal harmonics* to various familiar combinatorial enumerations; in particular we prove that the dimension of the space of diagonal harmonics is $(n+1)^{n-1}$. In order to prove these results, we have to work out some new results about geometry of the *Hilbert scheme of points in the plane* and a certain related algebraic variety. As a technical tool for our geometric results, in turn, we need to do some commutative algebra, which although complicated, has a quite explicit and combinatorial nature.

Today I want to give a short history of the original problem and how I got mixed up in it, and then show you some theorems in three seemingly unrelated areas. My goal in the remaining lectures will be to explain what these theorems have to do with each other and indicate how they are proved. The basic references are the series of four papers [15], [16], [17], and [18] (an abbreviated version of the last one appeared as [19]).

* Research supported in part by NSF grant DMS-0070772.

1.2. HISTORY

In 1988, Macdonald created something of a revolution in the ancient and classical theory of symmetric functions with the introduction of *Macdonald polynomials*. They are symmetric functions $P_\mu(x;q,t)$ in variables $x = x_1, x_2, \ldots$, with coefficients that are rational functions of two parameters q and t. Their importance stems in part from the fact that by specializing the two parameters in different ways we recover two previously known and important families of symmetric functions involving one parameter: the *Hall-Littlewood polynomials* (by setting $q = 0$) and the *Jack polynomials* (by setting $t = q^\alpha$ and letting $q \to 1$). After suitably normalizing and transforming the polynomials P_μ, we get polynomials $\tilde{H}_\mu(x;q,t)$, whose expansions in terms of Schur functions we may write as

$$\tilde{H}_\mu(x;q,t) = \sum_\lambda \tilde{K}_{\lambda,\mu}(q,t) s_\lambda(x).$$

Here λ and μ are partitions of an integer n. The coefficients $\tilde{K}_{\lambda,\mu}(q,t)$ are called Kostka-Macdonald coefficients. We'll define $\tilde{H}_\mu$ later. Right now we only mention that based on hand calculations, Macdonald conjectured that the Kostka-Macdonald coefficients are *polynomials with non-negative integer coefficients*:

$$\tilde{K}_{\lambda,\mu}(q,t) \in \mathbb{N}[q,t].$$

This is more remarkable for the fact that as defined, the $\tilde{K}_{\lambda,\mu}(q,t)$ are rational functions of q and t, and were only proved to be polynomials around 1996, in five independent papers by a total of seven authors [12, 13, 21, 22, 25].

Macdonald defined the coefficients $\tilde{K}_{\lambda,\mu}(q,t)$ in such a way that on setting $q = 0$ (the specialization from Macdonald to Hall-Littlewood polynomials) they yield the famous t-Kostka coefficients $\tilde{K}_{\lambda,\mu}(t) = \tilde{K}_{\lambda,\mu}(0,t)$. These were known to be in $\mathbb{N}[t]$ as a result of a cohomological interpretation due to Hotta and Springer [20, 26], and later a beautiful and subtle combinatorial interpretation due to Lascoux and Schützenberger [23]. Both descriptions are rather difficult.

When I came to U.C. San Diego in 1991, Adriano Garsia and Claudio Procesi had been working on a simpler approach to the positivity theorem for the t-Kostka coefficients [11], with the idea that it might extend to the q,t case. Adriano and I soon found the right extension—or so we conjectured—but surprisingly, our conjecture defied every attempt at an elementary proof. All that resulted from our attempts was an ever-larger pile of conjectures, notably those on diagonal harmonics alluded to above. Later, we discussed our efforts with Procesi. He was familiar with the geometry of the Hilbert scheme of points in the plane, and realized that there was a way it might explain the diagonal harmonics conjectures. What I'll describe in these lectures is ultimately the result of following up on this suggestion by Procesi. His proposed set-up eventually turned out to explain the diagonal harmonics conjectures and, still better and unexpectedly, to explain our original conjecture on the Kostka-Macdonald coefficients too.

Now let's turn to our promised list of seemingly unrelated theorems.

1.3. COMBINATORICS/LINEAR ALGEBRA

Let $M \subseteq \mathbb{N} \times \mathbb{N}$ be a finite subset of the first quadrant integer lattice, with $|M| = n$ elements. Often, M will be the Young diagram of a partition μ of n:

$$d(\mu) = \{(p,q) : p < \mu_{q+1}\}.$$

For example, the diagram of $\mu = (2,2,1)$ looks like this:

$$\begin{array}{ll} \bullet\,(2,0) & \\ \bullet\,(1,0) & \bullet\,(1,1) \\ \bullet\,(0,0) & \bullet\,(0,1). \end{array}$$

Note that origin of indices is $(0,0)$, and that the first coordinate, which indexes the rows of the diagram, is the coordinate along the vertical axis.

Given M, we define a polynomial $\Delta_M(x_1, y_1, \ldots, x_n, y_n)$ in $2n$ variables to be the $n \times n$ matrix determinant

$$\Delta_M(\mathbf{x}, \mathbf{y}) = \det \left[x_i^{p_j} y_i^{q_j} \right]_{1 \le i,j \le n}, \tag{1}$$

where

$$M = \{(p_1, q_1), \ldots, (p_n, q_n)\}.$$

When M is the diagram of a partition we abbreviate $\Delta_{d(\mu)}$ to Δ_μ. For example, for the partition $\mu = (2,2,1)$ whose diagram is displayed above, we have

$$\Delta_\mu = \det \begin{bmatrix} 1 & y_1 & x_1 & x_1 y_1 & x_1^2 \\ 1 & y_2 & x_2 & x_2 y_2 & x_2^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & y_5 & x_5 & x_5 y_5 & x_5^2 \end{bmatrix}.$$

Observe that Δ_μ is doubly homogeneous of degree $n(\mu) = \sum_i (i-1)\mu_i = \sum_i p_i$ in $\mathbf{x}$ and degree $n(\mu')$ in $\mathbf{y}$, where μ' denotes the conjugate partition. Observe also that there is a symmetry

$$\Delta_{\mu'}(\mathbf{x}, \mathbf{y}) = \Delta_\mu(\mathbf{y}, \mathbf{x}).$$

We let the symmetric group S_n act diagonally on the polynomial ring $\mathbb{C}[\mathbf{x}, \mathbf{y}]$, that is, by the rule

$$w x_i = x_{w(i)}, \quad w y_i = y_{w(i)} \qquad \text{for all } w \in S_n.$$

Then S_n permutes the rows of the matrix in (1), hence acts on Δ_M by

$$w\Delta_M = \varepsilon(w)\Delta_M,$$

where ε is the sign character. In other words, Δ_M is an *alternating* polynomial. Given any monomial $\mathbf{x}^p\mathbf{y}^q = x_1^{p_1}y_1^{q_1}\cdots x_n^{p_n}y_n^{q_n}$ with distinct exponents (p_i, q_i), the alternation of $\mathbf{x}^p\mathbf{y}^q$ is $\pm\Delta_M$ for the corresponding M. If the exponents are not distinct, then the alternation of the monomial $\mathbf{x}^p\mathbf{y}^q$ is zero. From this it is not hard to see that the set of all determinants Δ_M is a basis of the space of alternating polynomials $\mathbb{C}[\mathbf{x},\mathbf{y}]^\varepsilon$. Another way to see this is by identifying $\mathbb{C}[\mathbf{x},\mathbf{y}]^\varepsilon$ with the exterior power $\wedge^n\mathbb{C}[x,y]$. Then the determinants Δ_M are identified with the basis given by wedges of monomials $x^py^q \in \mathbb{C}[x,y]$.

Now, given a partition μ of n, consider the space spanned by all iterated partial derivatives of Δ_μ

$$D_\mu = \mathbb{C}[\partial\mathbf{x}, \partial\mathbf{y}]\Delta_\mu.$$

This space is

- finite dimensional,
- closed under differentiation (*i.e.*, it's a *Macaulay inverse system*),
- S_n-invariant, and
- doubly graded: $D_\mu = \bigoplus_{r,s}(D_\mu)_{r,s}$,

where $(D_\mu)_{r,s}$ is the subset of doubly homogeneous polynomials in D_μ of degree r in $\mathbf{x}$ and s in $\mathbf{y}$. The S_n action on D_μ respects the double grading.

THEOREM 1.1. *We have* $\dim D_\mu = n!$*, and* S_n *acts on* D_μ *by the regular representation.*

A refinement of this theorem describes the S_n action on each doubly homogeneous component $(D_\mu)_{r,s}$ individually. Recall that every S_n-module is a direct sum of irreducible ones, and that the irreducible S_n-modules V^λ (up to isomorphism) are indexed by partitions λ of n. The character of an S_n-module will be denoted $\mathrm{ch}V$. The irreducible characters are $\chi^\lambda = \mathrm{ch}V^\lambda$. The multiplicity of χ^λ in an arbitrary character ϕ is denoted $\langle\chi^\lambda, \phi\rangle$.

THEOREM 1.2. *The generating function for the multiplicty of* χ^λ *in the components* $(D_\mu)_{r,s}$ *is given by*

$$\sum_{r,s} t^r q^s \langle\chi^\lambda, \mathrm{ch}(D_\mu)_{r,s}\rangle = \tilde{K}_{\lambda,\mu}(q,t),$$

the Kostka-Macdonald coefficient.

The character multiplicities in the above formula are of course nonnegative integers, so this proves the Macdonald positivity conjecture.

COROLLARY 1.3. (Conjecture, Macdonald 1988). *We have* $\tilde{K}_{\lambda,\mu}(q,t) \in \mathbb{N}[q,t]$.

Example (the classical case). Take $\mu = (1^n)$. The diagram of (1^n) is

$$\begin{array}{l} \bullet\ (n-1,0) \\ \vdots \\ \bullet\ (1,0) \\ \bullet\ (0,0) \end{array}$$

Notice that the y variables have exponent zero in the determinant $\Delta_{(1^n)}$, while the x variables form the Vandermonde matrix. Thus we have

$$\Delta_{(1^n)} = \Delta(\mathbf{x}) = \prod_{i<j}(x_i - x_j),$$

the Vandermonde determinant. Now let

$$I = \{f(\mathbf{x}) : f(\partial\mathbf{x})\Delta(\mathbf{x}) = 0\}$$

be the annihilating ideal of the inverse system generated by $\Delta(\mathbf{x})$. It is equal to the ideal generated by all S_n-invariant polynomials in $\mathbf{x}$ without constant term, or just by the power-sums $p_k = \sum_{i=1}^n x_i^k$, that is,

$$I = (\mathbb{C}[\mathbf{x}]_+^{S_n}) = (p_1, \ldots, p_n).$$

We have isomorphisms of graded S_n-modules

$$D_{(1^n)} \cong \mathbb{C}[\mathbf{x}]/I \cong H^*(\mathrm{GL}_n/B).$$

On the far right here we have the cohomology ring of the flag variety for GL_n, which is well-known to have dimension $n!$ and to carry the regular representation of S_n. A corresponding result holds for any semi-simple complex Lie group and its Weyl group.

The spaces D_μ provide one kind of bivariate, doubly graded generalization of the above classical example. Another way to get such a generalization is to consider the Macaulay inverse system defined by the bivariate analog of the ideal I. Thus we set

$$J = (\mathbb{C}[\mathbf{x},\mathbf{y}]_+^{S_n}) = (p_{h,k} : 1 \leq h+k \leq n).$$

Here $p_{h,k} = \sum_{i=1}^n x_i^h y_i^k$ is a *polarized* power-sum; they generate J by a theorem of Weyl. The ideal J defines a Macaulay inverse system

$$DH_n = \{q \in \mathbb{C}[\mathbf{x},\mathbf{y}] : f(\partial\mathbf{x},\partial\mathbf{y})q = 0 \text{ for all } f \in J\}.$$

Or, more simply, DH_n is the solution space of the system of differential equations $p_{h,k}(\partial\mathbf{x}, \partial\mathbf{y})q = 0$. We have

$$DH_n \cong \mathbb{C}[\mathbf{x}, \mathbf{y}]/J$$

as (doubly) graded S_n-modules, so for character considerations the two are interchangable. The ring on the right is the ring of *coinvariants* for the diagonal action of S_n. The space DH_n is the space of harmonics for this action, or *diagonal harmonics.*

THEOREM 1.4. *We have* $\dim DH_n = (n+1)^{n-1}$.

This theorem also has a refinement giving the character of $(DH_n)_{r,s}$ in every degree r, s. The formula involves Macdonald polynomials and will be presented later. Garsia and I conjectured the above theorem around 1992. Shorty afterwards we and others including Gessel and Stanley found various combinatorial refinements to the conjecture. For example, $(n+1)^{n-1}$ counts *rooted forests* on the vertex set $\{1, \dots, n\}$. If we only look at the grading of DH_n by $\mathbf{x}$-degree, we find that $\dim(DH_n)_{d,-}$ is the number of rooted forests with d inversions (a pair of vertices $i < j$ is an *inversion* if j is on the path connecting i to the root of its tree). As another example, the dimension of the space DH_n^ε of alternating diagonal harmonics turns out to be the Catalan number C_n. A full discussion of these various conjectures can be found in [14].

Armed with Procesi's idea about the underlying geometry, I was eventually able to conjecture the full character formula for DH_n. Garsia and I then succeeded in proving that all the combinatorial conjectures would follow from the master formula [10]. The proof is based on the known specializations of Macdonald polynomials for $q = 1$ and for $q = t$, plus a lot of work with symmetric function identities and Garsia's theory of q-Lagrange inversion. This is very beautiful, but could take several lectures in itself, so I won't delve further into it.

1.4. ALGEBRAIC GEOMETRY

We denote by

$$H_n = \mathrm{Hilb}^n(\mathbb{C}^2)$$

the *Hilbert scheme of points in the affine plane* $\mathbb{C}^2$. It is an algebraic variety which parametrizes finite subschemes of length n in $\mathbb{C}^2$. You don't really need to know what those words mean, because (by definition) such subschemes correspond one-to-one with certain ideals in the coordinate ring of $\mathbb{C}^2$, and I'm going to tell you which ones they are. The coordinate ring of $\mathbb{C}^2$ is of course the ring $R = \mathbb{C}[x, y]$ of polynomials in the coordinate functions x and y. We can describe the Hilbert scheme set-theoretically as

$$H_n = \{I \subseteq R : \dim_{\mathbb{C}} R/I = n\}.$$

Thus the relevant ideals are those for which R/I is a finite-dimensional vector space, of dimension n. In more geometric language, R/I has Krull dimension zero, and *length* n. For such an ideal I, the locus $V(I) \subseteq \mathbb{C}^2$ defined by the vanishing of all polynomials in I is finite, with at most n points. But $V(I)$ may have fewer than n points, and then the local rings $(R/I)_P$ at these points may be non-reduced (they may contain non-zero nilpotent elements) and have length greater than 1. Defining the *multiplicity* of R/I at P to be the length of the local ring $(R/I)_P$, the sum of the multiplicities is always equal to n. The multiplicities do not determine I in general, so a point of H_n carries with it more information than just the set $V(I)$ and the multiplicities of its elements.

Examples. (1) The "generic" example of a point of H_n is the ideal $I = I(S)$ of all polynomials which vanish on a specified finite subset $S \subseteq \mathbb{C}^2$ of size $|S| = n$. In this case, I is a radical ideal and R/I is reduced. Each point in S has multiplicity one, and R/I can be identified with the ring of polynomial functions on S. Since S is finite, *every* function is a polynomial function, so $R/I \cong \mathbb{C}^S \cong \mathbb{C}^n$, showing that $I(S)$ is a legitimate element of H_n.

(2) The "most special" examples of points of H_n are the monomial ideals. Let μ be a partition of n, and let

$$I_\mu = \mathbb{C} \cdot \{x^p y^q : (p, q) \notin d(\mu)\}$$

be the ideal generated, and in fact spanned, by the monomials whose exponents (p, q) lie *outside* the diagram of μ. For example, we have

$$I_{(3,1)} = (x^2, xy, y^3).$$

Then the monomials $x^p y^q$ with (p, q) *inside* the diagram form a basis of R/I, and since there are n of them, we have $I_\mu \in H_n$.

Note that we always have $(x^n, y^n) \subseteq I_\mu$, and therefore $V(I_\mu) = \{\underline{0}\}$ as a set. The only point is the origin, and its multiplicity is n. Here is the way you should think of I_μ geometrically. The diagram of μ is an n-element subset $S = d(\mu)$ of $\mathbb{C}^2$ and has an ideal $I = I(S) \in H_n$, as in example (1) above. For any non-zero complex parameter u, the same goes for the rescaled set uS, and we have a family of ideals $I_u = I(uS) \in H_n$. They form a curve in H_n, and I_μ is their limit as $u \to 0$. In other words, I_μ and all the ideals I_u lie on a single algebraic curve $\mathbb{C}^1 \hookrightarrow H_n$, with I_μ corresponding to $u = 0$. Thus I_μ should be thought of as the ideal of an "infinitesimal" copy of the diagram $d(\mu)$ concentrated at the origin.

We haven't yet said anything about the structure of H_n as an algebraic variety, that is, about how it acquires local coordinates. Roughly speaking, its coordinates are defined similarly to the Plücker coordinates on a Grassmann variety. This will be explained in more detail later. For now, let's accept that it has a natural variety structure. It is a very special variety, as indicated by the following famous and remarkable theorem.

THEOREM 1.5. (Fogarty [7]). *The Hilbert scheme H_n is non-singular and irreducible, of dimension $2n$.*

This theorem does *not* hold for the Hilbert scheme of points in $\mathbb{C}^d$ when $d > 2$. We may mention here also that the "$n!$" theorem, Theorem 1.1, does not hold in more than two sets of variables $\mathbf{x}, \mathbf{y}, \ldots, \mathbf{z}$, and these two phenomena are connected with each other. One reason why elementary attempts to prove the $n!$ theorem have failed so far is that most ideas one thinks of are equally applicable to three or more sets of variables—and therefore must be wrong. In the geometric proof, the critical ingredient that causes it to break down beyond the bivariate case is the role of H_n.

We remark that the irreducibility aspect of Forgarty's theorem means that the ideals $I = I(S)$ in example (1) above are truly "generic," in that they form a dense (open) subset of H_n. Thus every $I \in H_n$ can be realized as a limit of ideals $I(S)$, somewhat as we did above for I_μ (but not always by a rescaling; that only gives the homogeneous ideals).

There is a map

$$\sigma \colon H_n \to S^n\mathbb{C}^2,$$

where $S^n\mathbb{C}^2 = \mathbb{C}^{2n}/S_n$ is the variety of *unordered* n-tuples $[\![P_1, \ldots, P_n]\!]$ of points in $\mathbb{C}^2$. It is defined by

$$\sigma(I) = [\![m_1 \cdot P_1, \ldots, m_k \cdot P_k]\!],$$

where $V(I) = \{P_1, \ldots, P_k\}$ and $m_i = \mathrm{length}(R/I)_{P_i}$ is the multiplicity of P_i. The map σ is called the *Chow morphism*, and is a morphism of algebraic varieties (in fact, a projective morphism). Note that for $[\![P_1, \ldots, P_n]\!] \in S^n\mathbb{C}^2$ with all P_i distinct, there is a unique $I \in H_n$ such that $\sigma(I) = [\![P_1, \ldots, P_n]\!]$, namely $I = I(S)$, where $S = \{P_1, \ldots, P_n\}$. Thus the Chow morphism is generically one-to-one, or *birational*, and since H_n is non-singular it is a *resolution of singularities* of $S^n\mathbb{C}^2$. Later we will see that it can also be described as a certain *blowup* of $S^n\mathbb{C}^2$.

Now we are ready for one last definition and our next theorem.

DEFINITION 1.6. *The isospectral Hilbert scheme X_n is the reduced fiber product*

$$\begin{array}{ccc} X_n & \longrightarrow & \mathbb{C}^{2n} \\ \downarrow & & \downarrow \\ H_n & \xrightarrow{\sigma} & S^n\mathbb{C}^2. \end{array}$$

In other words, X_n is the closed subset $\{(I, P_1, \ldots, P_n) : \sigma(I) = [\![P_1, \ldots, P_n]\!]\}$ $\subseteq H_n \times \mathbb{C}^{2n}$, with the induced structure of reduced algebraic variety.

For the experts, we should point out that the *scheme-theoretic* fiber product in this diagram, the closed subscheme of $H_n \times \mathbb{C}^{2n}$ whose ideal sheaf is described by the equations $\sigma(I) = [\![P_1, \ldots, P_n]\!]$, is *not* reduced. By definition, the ideal sheaf of X_n is the radical of the former ideal sheaf. We don't know how to give explicit local generators of the ideal sheaf of X_n, although we will give an implicit description later on.

THEOREM 1.7. *The isospectral Hilbert scheme X_n is normal, Cohen-Macaulay, and Gorenstein.*

We'll discuss the terms *normal*, *Cohen-Macaulay* and *Gorenstein* later, when they are needed. For now, they just mean that X_n, although singular, has only very special singularities.

1.5. COMMUTATIVE ALGEBRA

Let E be any space (topological space, algebraic variety, whatever). We define the *polygraph* $Z(n, l)$ to be the following subset of $E^n \times E^l$:

$$Z(n, l) = \Big\{(P_1, \ldots, P_n, Q_1, \ldots, Q_l) : Q_i \in \{P_1, \ldots, P_n\} \text{ for all } i\Big\}.$$

The reason for the name is as follows. Given a function $f\colon \{1, \ldots, l\} \to \{1, \ldots, n\}$, define the map

$$\pi_f\colon E^n \to E^l$$
$$\pi_f(P_1, \ldots, P_n) = (P_{f(1)}, \ldots, P_{f(l)}).$$

Its graph is the locus $W_f \subseteq E^n \times E^l$ defined by

$$W_f = \{(P_1, \ldots, P_n, Q_1, \ldots, Q_l) :\ Q_i = P_{f(i)} \text{ for all } i\}.$$

Clearly $Z(n, l) = \bigcup_f W_f$ is the union of these graphs, hence "polygraph."

Now fix $E = \mathbb{C}^2$ and fix coordinates on $E^n \times E^l = \mathbb{C}^{2n+2l}$

$$\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b} = x_1, y_1, \ldots, x_n, y_n, a_1, b_1, \ldots, a_l, b_l.$$

In coordinates, W_f is the linear subspace, and algebraic subvariety, defined by the equations

$$W_f = V(I_f), \quad I_f = (a_i - x_{f(i)}, b_i - y_{f(i)} : 1 \le i \le l).$$

Therefore $Z(n,l)$ is an arrangement of linear subpaces in $\mathbb{C}^{2n+2l}$, whose coordinate ring as a (non-irreducible) algebraic variety is

$$R(n,l) = \mathbb{C}[\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b}]/I(n,l),$$

where $I(n,l) = \bigcap_f I_f$ is the ideal of all polynomials vanishing on $Z(n,l)$ (or equivalently, vanishing on W_f for all f).

THEOREM 1.8. *For $E = \mathbb{C}^2$, the coordinate ring $R(n,l)$ of the polygraph $Z(n,l)$ is a free $\mathbb{C}[\mathbf{y}]$-module.*

By symmetry we could equally well have said that $R(n,l)$ is a free $\mathbb{C}[\mathbf{x}]$-module. The point is that it's free over the polynomial ring in either one of the two sets of coordinates on $\mathbb{C}^{2n}$.

1.6. CONCLUSION

We close with a hint as to the relationships between the various theorems discussed above, and an outline of the remaining lectures. We shall see that the $n!$ theorem is a consequence of the theorem on the geometry of X_n, Theorem 1.7. We can already indicate how the two are connected. Because H_n is nonsingular and the projection $\rho: X_n \to H_n$ is finite, the Cohen-Macaulay property of X_n is equivalent to ρ being *flat*. This means that its scheme-theoretic fibers have constant length. Now for a generic point $I = I(S)$ in H_n, with $S = \{P_1, \ldots, P_n\}$, the fiber of ρ over I consists of the points $(I, P_{w(1)}, \ldots, P_{w(n)})$, where $(P_{w(1)}, \ldots, P_{w(n)})$ is one of the $n!$ possible orderings of the points of S. These generic fibers can be identified with reduced, regular S_n-orbits in $\mathbb{C}^{2n}$, and have length $n!$. By flatness, every fiber has length $n!$, and carries the regular representation of S_n on its coordinate ring (the S_n-character is constant, as well as the length). Since $V(I_\mu) = \{\underline{0}\}$, the fiber of ρ over I_μ is (set-theoretically) concentrated at a unique point $Q_\mu = (I, \underline{0}, \ldots, \underline{0}) \in X_n$. The coordinate ring of the scheme-theoretic fiber is a non-reduced local ring of the form $R_\mu = \mathbb{C}[\mathbf{x}, \mathbf{y}]/J_\mu$, of length $n!$ and carrying the regular representation of S_n. It turns out that the ideal J_μ is exactly the annihilating ideal of the Macaulay inverse system D_μ, so that R_μ and D_μ are isomorphic as doubly-graded S_n-modules.

The proof of Theorem 1.7 is mostly geometric. Beginning with the next lecture, we will work out some basic descriptive facts about H_n, X_n and a *nested* Hilbert scheme $H_{n-1,n}$ to be introduced later, and then go through this geometric proof. At one point, however, we will assume a technical result: that the composite projection $X_n \to \mathbb{C}^{2n} \to \mathbb{C}^n$, where the second map is projection on the $\mathbf{y}$ coordinates, is flat. To justify this we need the polygraph theorem, Theorem 1.8. We will see that X_n is a blowup of $\mathbb{C}^{2n}$, and using Theorem 1.8, we will see that the Rees algebra defining the blowup is a free $\mathbb{C}[\mathbf{y}]$-module. This implies the required flatness result.

The proof of the polygraph theorem has a completely different flavor from the geometric argument, and we will come to it afterwards. In essence, it is simple: to prove that $R(n,l)$ is a free $\mathbb{C}[\mathbf{y}]$-module, we'll construct a free module basis. In practice, the inductive procedure for doing this is rather horrific. What's worse, the algorithm by which we construct the basis elements does not immediately show that they are in $R(n,l)$ at all! We will construct them as functions on $Z(n,l)$, but not obviously as regular functions, *i.e.*, functions defined by polynomials. So as an added complication we must prove as we go along that our basis elements are regular functions. We will not go through every detail of the proof in these lectures. What I will try to do is show you enough of the general framework and method of the construction so that you can more easily follow the details in [17], if you are interested.

What I will also try to show you, which is not in [17], is that the basis construction has some beautiful combinatorics. Specifically, the basis elements will be indexed by simple combinatorial data from which you can read off their degrees (like everything in this story, they are doubly homogeneous) and other identifying information. As a consequence we get a combinatorial interpretation of the doubly graded Hilbert series $A^{n,l}(q,t)$ of the ring $R(n,l)$. This has the remarkable feature that although the Hilbert series has an obvious symmetry $A^{n,l}(q,t) = A^{n,l}(t,q)$, the combinatorial description is utterly asymmetric. It also has the remarkable feature that there is a *formula* for $A^{n,l}(q,t)$ in terms of symmetric function operators derived from Macdonald polynomials. There are now many such quantities for which we can prove that the coefficients are non-negative by geometric means, but we only have combinatorial interpretations for two of them. One is $A^{n,l}(q,t)$, and the other is the q,t-Catalan number $C_n(q,t)$ whose combinatorial interpetation was discovered by Jim Haglund and proved by him and Garsia [8, 9]. Apart from the $n!$ theorem, we still have no purely combinatorial interpretation for $\tilde{K}_{\lambda,\mu}(q,t)$.

After all this, we will return to Theorem 1.4 and its refinement giving the full character formula for the diagonal harmonics. This is proved using the foregoing geometric results and something new, namely, cohomology vanishing theorems for tensor powers of the tautological bundle. Our main result on X_n can be interpreted (as I will explain) as an isomorphism between H_n and a *Hilbert scheme of regular orbits* of S_n acting on $\mathbb{C}^{2n}$. Given this isomorphism, we can apply an amazing recent theorem of Bridgeland, King and Reid to *completely characterize* the derived category of coherent sheaves on H_n. This has the tremendous virtue that it reduces the proofs of the vanishing theorems we need to "mere calculations." More exactly, it reduces them to calculations that can be done with the aid of the polygraph theorem! I intend to explain enough of this for you to get a taste of it, but in a sketchier way than the other material, since the sheaf cohomology techniques will be less accessible to many in the audience than the more intuitive geometry and commutative algebra involved in the study of Hilbert schemes and polygraphs.

2. Welcome to Hilbert schemes

In this lecture and the next we will explore how the properties of H_n and X_n are connected with the $n!$ theorem, and describe H_n and X_n in more elementary terms. We will also meet the *nested* Hilbert scheme $H_{n-1,n}$ which will play a key role in the geometric proof of Theorem 1.7 by induction on n.

2.1. HILBERT SCHEMES AS BLOWUPS

We recall some basics of algebraic geometry. Let R be a reduced (no non-zero nilpotent elements), finitely generated algebra over $\mathbb{C}$. If $x_1, \ldots, x_m$ generate R, we have

$$R = \mathbb{C}[\mathbf{x}]/I,$$

where I is an ideal such that $I = \sqrt{I}$. Then R is the ring of regular functions on the affine algebraic locus

$$V(I) \subseteq \mathbb{C}^n,$$

that is, its *coordinate ring*. The locus $V(I)$ can be identified with the set of closed points of a scheme, denoted $\operatorname{Spec} R$. You can and should think of $\operatorname{Spec} R$ as just another name for the locus $V(I)$—in the present context (reduced schemes of finite type over $\mathbb{C}$) the two concepts are essentially interchangeable.

A general reduced scheme X of finite type over $\mathbb{C}$ is a space which can be covered by finitely many affine sets $U = \operatorname{Spec} R$ for various rings R of the above kind. The notion of "regular function" on X is defined locally, that is, the regular functions form a *sheaf* of rings of functions on X, denoted $\mathcal{O}_X$. Thus $\mathcal{O}_X(U)$ is the ring of regular functions on the open set U. If U is an affine open set $U = \operatorname{Spec} R$, then $\mathcal{O}_X(U) = R$. The ring $\mathcal{O}_X(X)$ is the ring of *global* regular functions defined on all of X. If X is not affine, it may have few global regular functions: for example, the only global regular functions on projective space $\mathbb{P}^n$ are the constants.

Next let

$$S = S_0 \oplus S_1 \oplus \cdots$$

be a graded algebra, generated by S_0 and S_1. For $f \in S_1$ define

$$R_f = S[f^{-1}]_0 = S_0[f^{-1}S_1] \subseteq S[f^{-1}].$$

The affine schemes $U_f = \operatorname{Spec} R_f$ fit together to cover a scheme X projective over $\operatorname{Spec} S_0$. This scheme is denoted $\operatorname{Proj} S$. For each f, the canonical ring homomorphism $S_0 \to R_f$ induces a morphism of affine schemes $U_f \to \operatorname{Spec} S_0$. We have a canonical projective morphism $X \to \operatorname{Spec} S_0$, given on each set U_f by these morphisms.

Examples. (1) Take $S = \mathbb{C}[x_0, \ldots, x_d]$ to be a polynomial ring with its usual grading. We have

$$R_{x_i} = \mathbb{C}[x_0/x_i, \ldots, x_d/x_i] \quad (x_i/x_i = 1 \text{ omitted}),$$

and $U_{x_i} = \operatorname{Spec} R_{x_i} \cong \mathbb{C}^d$ is the affine d-space with coordinates x_0/x_i, …, x_d/x_i. In this case

$$\operatorname{Proj} S = \mathbb{P}^d$$

is projective space, and the morphism

$$\operatorname{Proj} S \to \operatorname{Spec} S_0 = \operatorname{Spec} \mathbb{C}$$

is the trivial morphism from $\mathbb{P}^d$ to a point. A point of $\mathbb{P}^d$ given in projective coordinates as $p = (x_0 : x_1 : \ldots : x_d)$ belongs to the open set U_{x_i} iff $x_i \neq 0$. Then rescaling to make $x_i = 1$, we can write $p = (x_0/x_i : \ldots : 1 : \ldots : x_d/x_i)$, which exhibits $x_0/x_i, \ldots, x_d/x_i$ as coordinates on U_{x_i}.

(2) If R is a ring and $J \subseteq R$ is an ideal, take

$$S = R \oplus J \oplus J^2 \oplus \cdots,$$

the *Rees algebra* of J. We can identify S with the subring $R[tJ]$ of the polynomial ring $R[t]$ in an indeterminate t. The morphism

$$\pi\colon X \to Y, \qquad X = \operatorname{Proj} S, \quad Y = \operatorname{Spec} R$$

is the *blowup of* Y *along the subscheme* $Z = V(J)$, or the blowup *at the ideal* J. Among its properties:

- The ideal sheaf $\mathcal{I} \subseteq \mathcal{O}_X$ of the subscheme $\pi^{-1}(Z) \subseteq X$ is locally free on one generator. Locally on U_f, for $f \in S_1 = J$, it is given by the ideal $\mathcal{I}(U_f) = (f) \subseteq R_f$. Note that $\mathcal{I}(U_f)$ is by definition the ideal generated by the image of J in R_f, but R_f contains an element g/f for all $g \in J$, so f generates JR_f.

– Over the complement $W = Y \setminus Z$ of Z, the blowup restricts to an isomorphism $\pi^{-1}(W) \cong W$.

We are going to construct the Hilbert scheme H_n and the isospectral Hilbert scheme X_n as blowups. To this end we set

$$A = \mathbb{C}[\mathbf{x}, \mathbf{y}]^{\epsilon},$$

the space of alternating polynomials for the diagonal S_n action on $\mathbb{C}[\mathbf{x}, \mathbf{y}] = \mathbb{C}[x_1, y_1, \ldots, x_n, y_n]$, and let

$$J = \mathbb{C}[\mathbf{x}, \mathbf{y}]A$$

be the ideal generated by A. Recall that the polynomials Δ_M for all $M \subseteq \mathbb{N} \times \mathbb{N}$, $|M| = n$ form a basis of A, and hence generate the ideal J. We define spaces A^d in the obvious way for $d > 0$, and set $A^0 = \mathbb{C}[\mathbf{x}, \mathbf{y}]^{S_n}$, to ensure that we have

$$A^j A^k \subseteq A^{j+k},$$

even when j or k is zero. Note that we also have

$$J^d = \mathbb{C}[\mathbf{x}, \mathbf{y}]A^d.$$

THEOREM 2.1. *We have*

$$H_n \cong \operatorname{Proj} S,$$

where

$$S = A^0 \oplus A^1 \oplus A^2 \oplus \cdots,$$

as a scheme over $\operatorname{Spec} A^0 = S^n\mathbb{C}^2$, *that is, the canonical projective morphism for the* Proj *is the Chow morphism* $\sigma\colon H_n \to S^n\mathbb{C}^2$. *We also have*

$$X_n \cong \operatorname{Proj} \mathbb{C}[\mathbf{x}, \mathbf{y}][tJ],$$

the blowup of $\mathbb{C}^{2n}$ *at* J.

Remarks. (1) For d even, A^d is contained in the ring of invariants $A^0 = \mathbb{C}[\mathbf{x}, \mathbf{y}]^{S_n}$ as an ideal, and

$$S^{(2)} = A^0 \oplus A^2 \oplus A^4 \oplus \cdots$$

is the Rees algebra of A^2. A general fact about the Proj construction is that $\operatorname{Proj} S^{(k)} \cong \operatorname{Proj} S$, so H_n is the blowup of $S^n\mathbb{C}^2$ at the ideal A^2. The description as $\operatorname{Proj} S$ is preferable to the description as the blowup $\operatorname{Proj} S^{(2)}$ because it gives rise to the "correct" ample line bundle $\mathcal{O}(1)$ on H_n.

(2) The first part of Theorem 2.1 is proved in [15], and the part about X_n in [17]. In the latter paper I also prove that J is equal to its radical,

which implies the same for $A^2 = J \cap \mathbb{C}[\mathbf{x}, \mathbf{y}]^{S_n}$. In fact J is the ideal of the union of pairwise diagonals

$$\bigcup_{i \neq j} E_{ij} \subseteq \mathbb{C}^{2n}, \quad E_{ij} = V(x_i - x_j, y_i - y_j),$$

and A^2 is the ideal of its image in $S^n\mathbb{C}^2$, the set of points $[\![P_1, \ldots, P_n]\!]$ with the P_i not all distinct. Thus H_n and X_n are the blowups of $S^n\mathbb{C}^2$ and $\mathbb{C}^{2n}$ respectively, along the pairwise diagonals. This had been expected by geometers but not proved before.

Example. For $n = 2$, we have $J = (x_1 - x_2, y_1 - y_2)$, and these two generators also generate A as an A^0-module. The blowup $\operatorname{Proj} S$ is therefore covered by two affines

$$U_{x_1 - x_2} = \operatorname{Spec} A^0[\frac{y_1 - y_2}{x_1 - x_2}], \quad U_{y_1 - y_2} = \operatorname{Spec} A^0[\frac{x_1 - x_2}{y_1 - y_2}].$$

To see how this identifies with the Hilbert scheme H_2, observe that the latter is covered by two open sets

$$\begin{aligned} W_x &= \{I = (x^2 - e_1 x + e_2,\ y - a_1 x - a_0)\}, \\ W_y &= \{I = (y^2 - e_1' y + e_2',\ x - a_1' y - a_0')\}. \end{aligned}$$

Since the Chow morphism maps H_2 birationally on $S^2\mathbb{C}^2$, all regular functions on open sets in H_2 can be identified with S_2-invariant rational functions of x_1, x_2, y_1, y_2. For the coordinates e_1, e_2, a_0, a_1 on W_x we have

$$\begin{aligned} e_1 &= x_1 + x_2 = \tfrac{x_1^2 - x_2^2}{x_1 - x_2} \\ e_2 &= x_1 x_2 = \tfrac{x_1^2 x_2 - x_1 x_2^2}{x_1 - x_2} \\ a_0 &= \tfrac{x_1 y_2 - y_1 x_2}{x_1 - x_2} = \tfrac{1}{2}\left(y_1 + y_2 - \tfrac{y_1 - y_2}{x_1 - x_2}(x_1 + x_2)\right) \\ a_1 &= \tfrac{y_1 - y_2}{x_1 - x_2}. \end{aligned}$$

These equations can easily be verified when $I = I(S)$, where $S \subseteq \mathbb{C}^2$ is a set of two points with distinct x-coordinates $x_1 \neq x_2$. All the above expressions belong to the ring

$$A^0[\frac{y_1 - y_2}{x_1 - x_2}] = \mathbb{C}[x_1 + x_2, x_1 x_2, y_1 + y_2, \frac{y_1 - y_2}{x_1 - x_2}],$$

and they generate it. This gives an explicit isomorphism between W_x and $U_{x_1 - x_2}$, and there is a similar isomorphism between W_y and $U_{y_1 - y_2}$.

We will outline the proof of Theorem 2.1, but we first need to introduce the *universal family*

$$\begin{array}{c} F \subseteq H_n \times \mathbb{C}^2 \\ \pi \big\downarrow \qquad\qquad \\ H_n \qquad\qquad \end{array}$$

defined by

$$F = \{(I, P) : P \in V(I)\},$$

so that the fiber of the projection $\pi: F \to H_n$ over $I \in H_n$ is the subscheme $V(I) \subseteq \mathbb{C}^2$. In general, universal families over Hilbert schemes are not reduced, but it follows from Fogarty's theorem that F is reduced (F is flat and generically reduced over the reduced and irreducible scheme H_n). Hence the set-theoretic description of F above fully characterizes it. The scheme structure of H_n is defined by a universal property of the universal family F, whose details need not detain us here.

What we do need to explore is the way in which regular functions on F can be interpreted as sections of a sheaf on H_n. Specifically, we define

$$B = \pi_* \mathcal{O}_F$$

to be the sheaf whose set of sections $B(U)$ on any open set $U \subseteq H_n$ is the algebra of regular functions on the open set $\pi^{-1}(U) \subseteq F$. Note that any regular function $f \in \mathcal{O}_{H_n}(U)$ composes with π to give a regular function $\pi^* f = f \circ \pi$ on $\pi^{-1}(U)$. This gives a homomorphism of sheaves of algebras $\mathcal{O}_{H_n} \to \pi_* \mathcal{O}_F = B$, making B a sheaf of $\mathcal{O}_{H_n}$-algebras. These constructions of course make sense with any morphism of schemes in place of π. Since F is a closed subscheme of $H_n \times \mathbb{C}^2$, local coordinates on F are generated by coordinates pulled back from H_n, together with coordinates x, y on $\mathbb{C}^2$. This implies that B is generated as a sheaf of $\mathcal{O}_{H_n}$-algebras by x and y.

Our particular morphism π is *flat* and *finite*. Finiteness means that B is a *coherent* sheaf of $\mathcal{O}_{H_n}$-modules, that is, $B(U)$ is a finitely generated $\mathcal{O}_{H_n}(U)$-module. For a finite morphism, flatness means additionally that B is a *locally free* sheaf of $\mathcal{O}_{H_n}$-modules. In our case B is locally free of rank n, the common length of all scheme-theoretic fibers of π.

Example. If $e_1, \ldots, e_n$, $a_0, \ldots, a_{n-1}$ are arbitrary complex numbers, the ideal

$$I = (x^n - e_1 x^{n-1} + e_2 x^{n-2} - \cdots \pm e_n,\ y - a_{n-1} x^{n-1} - \cdots - a_1 x - a_0) \\ \subseteq \mathbb{C}[x, y] \subseteq R$$

has the property that $\{1, x, \ldots, x^{n-1}\}$ is a basis of R/I. These ideals form an open set $W_x \subseteq H_n$, isomorphic to $\mathbb{C}^{2n}$, with coordinates $e_1, \ldots, e_n$, $a_0, \ldots, a_{n-1}$. The coordinates on $W_x \times \mathbb{C}^2$ are $\mathbf{e}$, $\mathbf{a}$, x, y. The open set

$$\pi^{-1}(W_x) = F \cap (W_x \times \mathbb{C}^2)$$

in the universal family is defined by the two equations

$$x^n - e_1 x^{n-1} + e_2 x^{n-2} - \cdots \pm e_n = 0,$$

$$y - a_{n-1} x^{n-1} - \cdots - a_1 x - a_0 = 0$$

among these coordinates. We have $\mathcal{O}_{H_n}(W_x) = \mathbb{C}[\mathbf{e}, \mathbf{a}]$, and the algebra of sections $B(W_x)$ is the $\mathbb{C}[\mathbf{e}, \mathbf{a}]$-algebra $\mathbb{C}[\mathbf{e}, \mathbf{a}, x, y]/\mathcal{I}(W_x)$, where $\mathcal{I}(W_x)$ is generated by the two equations above. As a $\mathbb{C}[\mathbf{e}, \mathbf{a}]$-module, $B(W_x)$ is free with basis $\{1, x, \ldots, x^{n-1}\}$. More generally, given any set M of n monomials in x and y, if $W \subseteq H_n$ is the open set whose points are ideals I such that M is a basis of R/I, then $B(W)$ is a free $\mathcal{O}_{H_n}(W)$-module with basis M. Since H_n is covered by such sets W, this exhibits B as a locally free sheaf of rank n.

Now if V is an algebraic vector bundle of rank n over any scheme X, the sections of V form a sheaf on X, and this sheaf is locally free of rank n. Any vector space basis of the fiber $V(x)$ over a point $x \in X$ is also a free $\mathcal{O}_X(U)$-module basis of $V(U)$, for some open neighborhood U containing x. Conversely, every locally free sheaf V of rank n can be realized as the sheaf of sections of an algebraic vector bundle. In precise language, the fiber of this vector bundle at x is $V \otimes_{\mathcal{O}_X} k_x$, where $k_x = \mathcal{O}_{X,x}/x$ is the residue field of the local ring of X at the point x. We will abuse notation and write V both for the vector bundle and its sheaf of sections.

Subjecting the locally free sheaf B to these considerations identifies it with the *tautological bundle* on H_n whose fiber at I is the n dimensional vector space R/I. As such, B is a quotient bundle of the trivial bundle $\mathbb{C}[x, y] \otimes_{\mathbb{C}} \mathcal{O}_{H_n}$ with constant fiber $R = \mathbb{C}[x, y]$.

We now turn to the proof of Theorem 2.1. Given $M \subseteq \mathbb{N} \times \mathbb{N}$ with $|M| = n$, set

$$\mathcal{B}_M = \{x^p y^q : (p, q) \in M\},$$

and let $W_M \subseteq H_n$ be the open affine subset

$$W_M = \{I \in H_n : \mathcal{B}_M \text{ is a basis of } R/I\}, \tag{2}$$

so that $B(W_M)$ is free with basis $\mathcal{B}_M$. Since $B(I) = R/I$, the fiber of the tensor power $B^{\otimes n}$ at a point $I \in H_n$ can be identified with

$$\mathbb{C}[\mathbf{x}, \mathbf{y}]/(I(x_1, y_1) + \cdots + I(x_n, y_n)),$$

and the fiber of the exterior power $\wedge^n B$ can be identified with the image of $A = \mathbb{C}[\mathbf{x}, \mathbf{y}]^\epsilon$ in the above space. In particular, every $\Delta_L \in A$ defines a global section of $\wedge^n B$, and the section Δ_M is nowhere-vanishing on W_M. Now $\wedge^n B$ is a vector bundle of rank 1, or *line bundle*, so the ratio $\Delta_L/\Delta_M \in \mathcal{O}_{H_n}(W_M)$ is a well-defined regular function on W_M, for every Δ_L.

At a point $I = I(S) \in W_M$, where $S = \{P_1, \ldots, P_n\}$ for n distinct points $P_i \in \mathbb{C}^2$, the interpretation of Δ_L/Δ_M as a ratio of sections in $B(W_M)$ coincides with its interpretation as a rational function of $\mathbf{x}, \mathbf{y}$. Since such ideals $I(S)$ form a dense set, all identities among the functions Δ_L/Δ_M which hold in the ring of rational functions

$$\mathbb{C}[\Delta_L/\Delta_M : \text{all } L]$$

also hold in $\mathcal{O}(W_M)$, giving rise to a ring homomorphism

$$\mathbb{C}[\Delta_L/\Delta_M : \text{all } L] \to \mathcal{O}(W_M) \tag{3}$$

and corresponding morphism of schemes

$$W_M \to \operatorname{Spec} \mathbb{C}[\Delta_L/\Delta_M : \text{all } L] = U_{\Delta_M} \subseteq \operatorname{Proj} S.$$

Here $S = \bigoplus_d A^d$ is as in Theorem 2.1. These morphisms combine to give a morphism

$$\alpha \colon H_n \to \operatorname{Proj} S.$$

We are to show that α is an isomorphism.

Every ideal $I \in W_M$ is generated, and even spanned as a vector space, by elements of the form

$$x^h y^k - \sum_{(p,q)\in M} c_{pq}^{hk} x^p y^q. \tag{4}$$

The coefficients c_{pq}^{hk} are regular functions on W_M, since they give the section $x^h y^k \in B(W_M)$ in terms of the basis $\mathcal{B}_M$. Moreover, they generate the coordinate ring of W_M. The reason for this is that the universal family is defined over W_M by the vanishing of the polynomials in (4). One way of expressing the universal property of H_n and F, which defines H_n as a scheme, is that once we have enough coordinates to describe the ideal of the universal family, we have them all. A look back at the last example may clarify this somewhat: the coordinates $e_1, \ldots, e_n$ are the c_{pq}^{hk} for $(h,k) = (n,0)$, $x^h y^k = x^n$, while $a_0, \ldots, a_{n-1}$ are the c_{pq}^{hk} for $(h,k) = (0,1)$, $x^h y^k = y$.

For $I = I(S)$ with $S = \{P_1, \ldots, P_n\}$ as before, we must have

$$x_i^h y_i^k = \sum_{(p,q)\in M} c_{pq}^{hk} x_i^p y_i^q$$

for all $i = 1, \ldots, n$. Fixing h, k, these give n equations in n "unknowns" c^{hk}_{pq}, whose solution is given by Cramer's rule as

$$c^{hk}_{pq} = \Delta_L/\Delta_M, \quad L = M \cup \{(h,k)\} \setminus \{(p,q)\}.$$

Since the c^{hk}_{pq} generate $\mathcal{O}(W_M)$, this shows that the ring homomorphism in (3) is surjective, so α is a closed embedding. On the generic locus, where $I = I(S)$, both the Chow morphism and the morphism $\mathrm{Proj}\, S \to S^n\mathbb{C}^2$ restrict to isomorphisms, and hence so does α. Since the generic locus is dense, this implies that the closed embedding α is an isomorphism.

We have now proven that $H_n \cong \mathrm{Proj}\, S$, and want to prove that $X_n = \mathrm{Proj}\, T$, where $T = \mathbb{C}[\mathbf{x}, \mathbf{y}][tJ]$. Observe that S is a subring of T (in fact, S is the ring of invariants T^{S_n} if we agree to make S_n act on t by the sign character). Since $J^d = \mathbb{C}[\mathbf{x}, \mathbf{y}]A^d$, we see that $\mathbb{C}[\mathbf{x}, \mathbf{y}]$ and S together generate T, so $\mathrm{Proj}\, T$ is a closed subscheme of $\mathrm{Proj}\, S \times \mathrm{Spec}\, \mathbb{C}[\mathbf{x}, \mathbf{y}] = H_n \times \mathbb{C}^2$. We have a commutative diagram

$$\begin{array}{ccc} \mathrm{Proj}\, T & \longrightarrow & \mathbb{C}^{2n} \\ \downarrow & & \downarrow \\ H_n = \mathrm{Proj}\, S & \longrightarrow & S^n\mathbb{C}^2. \end{array}$$

Since $\mathrm{Proj}\, T$ is reduced, this implies $\mathrm{Proj}\, T \subseteq X_n$. A geometric argument using the irreducibility of H_n shows that X_n is also irreducible (see [17]). Since $\mathrm{Proj}\, T$ is closed in X_n and contains the generic locus, we have $\mathrm{Proj}\, T = X_n$. $\square$

Remarks. (1) For $d > 2$ the Hilbert scheme $\mathrm{Hilb}^n(\mathbb{C}^d)$ is not irreducible in general. It has a unique irreducible component which maps birationally on $S^n\mathbb{C}^d$, namely, the closure of the generic locus. The proof above adapts to show that this *generic component* is the blowup of $S^n\mathbb{C}^d$ at the ideal $A^2 \subseteq \mathbb{C}[\mathbf{x}, \mathbf{y}, \ldots, \mathbf{z}]^{S_n}$, where $A = \mathbb{C}[\mathbf{x}, \mathbf{y}, \ldots, \mathbf{z}]^{\epsilon}$ as in the $d = 2$ case. Similarly, the analog of X_n, the reduced fiber product of the generic component with $\mathbb{C}^{dn}$ over $S^n\mathbb{C}^d$, is the blowup of $\mathbb{C}^{dn}$ at the ideal $J = \mathbb{C}[\mathbf{x}, \mathbf{y}, \ldots, \mathbf{z}]A$. What we don't know is whether J and A^2 are equal to their radicals, that is, the ideals of the pairwise diagonals, for $d > 2$. We conjecture that this is true.

(2) Theorem 2.1 provides us with a compact, explicit description of the coordinate rings of affine open sets in H_n and X_n. Namely, the open set W_M in (2) has coordinate ring $\mathbb{C}[\Delta_L/\Delta_M : \text{all } L]$, and (as follows easily) the preimage of W_M by the projection $X_n \to H_n$ has coordinate ring $\mathbb{C}[\mathbf{x}, \mathbf{y}, \Delta_L/\Delta_M]$. To cover H_n it suffices to take M to be the diagram of a partition μ. Then the monomial ideal I_μ is the point of W_μ defined by the vanishing of all the coordinates Δ_L/Δ_μ, $L \neq \mu$. The scheme-theoretic

fiber of X_n over I_μ therefore has coordinate ring

$$\mathbb{C}[\mathbf{x}, \mathbf{y}, \Delta_L/\Delta_\mu]/(\Delta_L/\Delta_\mu), \tag{5}$$

that is, the quotient ring of $\mathbb{C}[\mathbf{x}, \mathbf{y}]$ obtained by adjoining the fractions Δ_L/Δ_μ and then modding them out. The content of Theorem 1.7 boils down to the statement that the ring in (5) is Gorenstein. Later we will see that it is none other than the ring

$$\mathbb{C}[\mathbf{x}, \mathbf{y}]/J_\mu, \tag{6}$$

where J_μ is the annihilating ideal of the Macaulay inverse system generated by Δ_μ. One might hope to give an elementary proof of Theorems 1.1 and 1.7 by showing directly that the two rings in (5) and (6) are the same, but so far this has not been done.

3. Towards the proof of the $n!$ theorem

3.1. AN EXAMPLE: $N = 3$

Let's take a look at the picture for $n = 3$ over a non-trivial affine open subset of the Hilbert scheme. First, what do points of H_3 look like, viewed as subschemes $S \subseteq \mathbb{C}^2$ of length 3? There are four kinds:

- $S =$ three reduced points.
- $S =$ a double point on a line L plus one reduced point. The scheme structure detects the slope of L.
- $S =$ a triple point on a curve C. The scheme structure detects the slope and curvature of C at the triple point.
- $S =$ a non-curvilinear triple point, for example $I = I_{(2,1)} = (x, y)^2$. This example is unique up to translation.

Now we focus attention on the open set

$$U_{(2,1)} = \left\{ I \in H_3 : \{1, x, y\} \text{ is a basis of } R/I \right\}.$$

In this example it happens that $U_{(2,1)}$ is the set of *non-collinear* subschemes S. In other words, a subscheme S as above is *not* in $U_{(2,1)}$ if S is either three collinear reduced points, or a double point and a reduced point on the same line, or a triple point on a line. We can also describe $U_{(2,1)}$ as the *basin of attraction* for the fixed point $I_{(2,1)}$ under the contraction $x \mapsto tx$, $y \mapsto ty$, as $t \to 0$. All non-collinear subschemes contract to $I_{(2,1)}$ in the limit (collinear ones contract to a triple point at $\underline{0}$ on a line through the origin).

Since the tautological bundle B has fibers $B(I) = R/I$, the open set $U_{(2,1)}$ is the non-vanising locus of the section $1 \wedge x \wedge y$ of the line bundle

$\wedge^3 B$. In terms of our blowup construction, this section is represented by the alternating polynomial $\Delta_{(2,1)}(\mathbf{x}, \mathbf{y})$. The algebra of regular functions on $U_{(2,1)}$ is therefore

$$\mathbb{C}[\Delta_L / \Delta_{(2,1)}].$$

Another way to describe the regular functions on $U_{(2,1)}$ is a follows. Every $I \in U_{(2,1)}$ has generators

$$x^2 - ax - by - g \tag{7}$$

$$xy - cx - dy - h \tag{8}$$

$$y^2 - ex - fy - j \tag{9}$$

for some complex parameters $a, b, \ldots, j$. Modulo the above generators we can reduce a monomial, say $x^2 y$, in more than one way:

$$x^2 y \to y(ax + by + g) \to \left\{ \begin{array}{c} a(cx + dy + h) \\ +b(ex + fy + j) \\ +gy \end{array} \right.$$

or

$$x^2 y \to x(cx + dy + h) \to \left\{ \begin{array}{c} c(ax + by + g) \\ +d(cx + dy + h) \\ +hx. \end{array} \right.$$

Equating these yields conditions on the parameters. Doing this for reductions of xy^2 as well, we find

$$\begin{aligned} h &= be - cd \\ g &= b(c - f) + d(d - a) \\ j &= e(d - a) + c(c - f). \end{aligned}$$

The remaining parameters serve as coordinates on $U_{(2,1)}$:

$$U_{(2,1)} = \operatorname{Spec} \mathbb{C}[a, b, \ldots, f] \cong \mathbb{C}^6.$$

We have the identification

$$\mathbb{C}[a, b, \ldots, f] = \mathbb{C}[\Delta_L / \Delta_{(2,1)}],$$

given by expressions for $a, b, \ldots, f$ such as

$$a = \Delta_L / \Delta_{(2,1)} \quad \text{for } L = \{(2,0), (0,0), (0,1)\}.$$

To get this expression for a, we note that a is the coefficient of x in x^2 (mod I), so that applying Cramer's rule, we obtain L as the diagram of

$\mu = (2,1)$ with the point $(1,0)$, corresponding to x, deleted and replaced by the point $(2,0)$, corresponding to x^2:

$$L = \begin{array}{cc} \bullet & \\ \circ & \\ \bullet & \bullet \end{array} \quad .$$

Now let's look at X_3, and specifically at the open set

$$\rho^{-1}(U_{(2,1)}) \quad \text{via} \quad \rho\colon X_3 \to H_3.$$

Its coordinate ring is $\mathbb{C}[\mathbf{x}, \mathbf{y}, \Delta_L/\Delta_{(2,1)}]$, or $\mathbb{C}[x_1, y_1, \ldots, x_3, y_3, a, b, \ldots, f]$ modulo some relations which are not so easy to determine except by brute computation. We claim that X_3 is locally Cohen-Macaulay and Gorenstein at the unique point

$$Q_{(2,1)} = (I_{(2,1)}, \underline{0}, \underline{0}, \underline{0})$$

lying above $I_{(2,1)}$. The coordinates $a, b, \ldots, f$ form a system of parameters, so the Cohen-Macaulay property is that the ring

$$\mathbb{C}[\mathbf{x}, \mathbf{y}, \Delta_L/\Delta_{(2,1)}] \tag{10}$$

is a free $\mathbb{C}[a, b, \ldots, f]$ module, necessarily of rank $6 = 3!$ because the finite morphism ρ has degree 6. To establish this, locally at $Q_{(2,1)}$ (and in this case also globally on $U_{(2,1)}$ because it is the basin of attraction for $I_{(2,1)}$), we look at the quotient ring

$$\begin{aligned} \mathbb{C}[\mathbf{x}, \mathbf{y}, \Delta_L/\Delta_{(2,1)}] \,/(\Delta_L/\Delta_{(2,1)}) \; &= \; \mathbb{C}[\mathbf{x}, \mathbf{y}, \Delta_L/\Delta_{(2,1)}]/(a, b, \ldots, f) \\ &= \; \mathbb{C}[\mathbf{x}, \mathbf{y}]/J \end{aligned}$$

for some ideal J. Geometrically, this is the coordinate ring of the scheme-theoretic fiber $\rho^{-1}(\{I_{(2,1)}\})$. It has Krull dimension zero and vector space dimension (or length) ≥ 6, with equality if and only if X_3 is locally Cohen-Macaulay. What are some elements of J? For one thing, we have $\mathbb{C}[\mathbf{x}, \mathbf{y}]_+^{S_n} \subseteq J$, for if $p \in \mathbb{C}[\mathbf{x}, \mathbf{y}]$ is S_n-invariant and homogeneous of degree > 0, then $\Delta_{(2,1)}p$ is in the linear span of $\{\Delta_L : L \neq \mu\}$. For another, we have x_i^2, $x_i y_i$, $y_i^2 \in J$ for $i = 1, 2, 3$ because the generators of I in (7) vanish on setting $x = x_i$, $y = y_i$, and all but their leading terms belong to $(a, b, \ldots, f)$. As it turns out, the full ideal J is equal to

$$J = J_\mu = \{p : p(\partial \mathbf{x}, \partial \mathbf{y}) \Delta_{(2,1)} = 0\}$$

for $\mu = (2,1)$. This can be verified with some effort by direct calculation of both J and J_μ, and shows all of the following:

– The $n!$ theorem holds for $\mu = (2,1)$, since $\dim \mathbb{C}[\mathbf{x}, \mathbf{y}]/J_\mu = 3! = 6$.

- X_3 is locally Cohen-Macaulay at $Q_{(2,1)}$.
- X_3 is *Gorenstein* at $Q_{(2,1)}$. The reason for this is that $\mathbb{C}[\mathbf{x},\mathbf{y}]/J_\mu$ is a Gorenstein Artin graded algebra. In fact a graded Artin algebra $\mathbb{C}[\mathbf{x}]/I$ is Gorenstein if and only if I is the annihilating ideal of a Macaulay inverse system generated by one element, which is the case for J_μ. Since $\mathbb{C}[\mathbf{x},\mathbf{y}]/J$ is the quotient of our Cohen-Macaulay ring in (10) by a parameter ideal, the Gorenstein property of the latter ring is equivalent to that of $\mathbb{C}[\mathbf{x},\mathbf{y}]/J$.

3.2. THE IDEAL SHEAF OF X_n

Our next order of business is to describe X_n as a closed subvariety of $H_n \times \mathbb{C}^{2n}$ by specifying its ideal sheaf, that is the subsheaf of $\mathcal{O}_{H_n\times\mathbb{C}^{2n}}$ consisting of regular functions that vanish on X_n. To do this we first observe that if $(I, P_1, \ldots, P_n)$ is a point of X_n, then each P_i belongs to $V(I)$, so (I, P_i) is a point of the universal family F. In other words, X_n is a subset of the n-fold fiber product

$$F^n/H_n = \{(I, P_1, \ldots, P_n) : P_i \in V(I) \text{ for all } i\} \subseteq H_n \times \mathbb{C}^{2n}.$$

Now we will describe the ideal sheaf of X_n as a subsheaf of the sheaf of regular functions on F^n/H_n. More precisely, we will describe its pushdown to H_n. We denote the projection of F onto H_n by π, as before, and (abusing notation slightly) we denote by ρ the projection onto H_n of $H_n \times \mathbb{C}^{2n}$, or of its subschemes X_n or F^n/H_n. The projections π and ρ are *affine*, meaning that the inverse image of any affine open subset $U \subseteq H_n$ is again affine. Hence $B = \pi_*\mathcal{O}_F$ contains full information about the sheaf of regular functions $\mathcal{O}_F$, and the same goes for $\rho_*\mathcal{O}_{F^n/H_n}$ and $\rho_*\mathcal{O}_{X_n}$. We have

$$B^{\otimes n} = \rho_*\mathcal{O}_{F^n/H_n}$$

and we define

$$P = B^{\otimes n}/\mathcal{J} = \rho_*\mathcal{O}_{X_n}.$$

Our problem is to determine the sheaf of ideals $\mathcal{J}$ in the sheaf of $\mathcal{O}_{H_n}$-algebras $B^{\otimes n}$. This problem is solved by the following result.

PROPOSITION 3.1. *Let*

$$\phi\colon B^{\otimes n} \to (B^{\otimes n})^* \otimes \wedge^n B$$

be the homomorphism of vector bundles (or of their sheaves of sections) induced by the pairing

$$B^{\otimes n} \otimes B^{\otimes n} \to \wedge^n B$$

which is multiplication in $B^{\otimes n}$ followed by alternation: $B^{\otimes n} \to \wedge^n B$. Then the ideal sheaf $\mathcal{J}$ of X_n is the kernel of ϕ:

$$\mathcal{J} = \ker \phi\,; \quad P \cong \operatorname{im} \phi.$$

Proof. Let

$$\Theta^\epsilon f = \sum_{w \in S_n} \epsilon(w)\; wf \tag{11}$$

be the alternation operator. By definition, a section g of $B^{\otimes n}$ belongs to $\ker \phi$ iff $\Theta^\epsilon(gs) = 0$ for all s. Note that the regular function on (an open set in) F^n/H_n represented by any alternating section of $B^{\otimes n}$ must vanish at every point $(I, P_1, \ldots, P_n)$ such that $P_i = P_j$ for some $i \neq j$. Furthermore, F^n/H_n is the union of the locus V consisting of such points, and X_n. Suppose g belongs to $\mathcal{J}$. Then so does $\Theta^\epsilon(gs)$, so the latter vanishes on both V and X_n. Therefore it vanishes identically as a section of $\wedge^n B$, and g belongs to $\ker \phi$. Conversely, suppose g does not belong to $\mathcal{J}$, so g does not vanish identically on X_n. We can find a point $x = (I, P_1, \ldots, P_n)$ in X_n with all P_i distinct and $g(x) \neq 0$, since the generic locus is dense in X_n. Multiplying g by a suitable s, we can arrange that $gs(wx) = 0$ for all $1 \neq w \in S_n$, but $gs(x) \neq 0$. Then $\Theta^\epsilon(gs)(x) \neq 0$, so $g \notin \ker \phi$. □

3.3. THE ANNIHILATING IDEAL OF Δ_μ

The significance of Proposition 3.1 is that its implicit description of the ideal sheaf $\mathcal{J}$ closely parallels a similar implicit description of the annihilating ideal $J_\mu \subseteq \mathbb{C}[\mathbf{x}, \mathbf{y}]$ for the Macaulay inverse system D_μ generated by Δ_μ.

PROPOSITION 3.2. *Let*

$$J_\mu = \{p \in \mathbb{C}[\mathbf{x}, \mathbf{y}] : p(\partial \mathbf{x}, \partial \mathbf{y}) \Delta_\mu = 0\}.$$

Then p belongs to J_μ if and only if the coefficient of Δ_μ in $\Theta^\epsilon(ps)$ is zero, for all $s \in \mathbb{C}[\mathbf{x}, \mathbf{y}]$.

Here Θ^ϵ is the alternation operator in (11). It makes sense to speak of the coefficient of Δ_μ in $\Theta^\epsilon(ps)$, because the polynomials Δ_M form a basis of $\mathbb{C}[\mathbf{x}, \mathbf{y}]$.

Proof. The coefficient $\Theta^\epsilon(ps)|_{\Delta_\mu}$ is, apart from a fixed scalar factor, just the constant term of $s(\partial \mathbf{x}, \partial \mathbf{y}) p(\partial \mathbf{x}, \partial \mathbf{y}) \Delta_\mu$. By Taylor's theorem, this vanishes for all s if and only if $p(\partial \mathbf{x}, \partial \mathbf{y}) \Delta_\mu = 0$. □

We can reformulate the above proposition in the following suggestive way. Recall that $I_\mu \in H_n$ denotes the ideal spanned by all monomials not in the diagram of μ. We see immediately that in the ring

$$B^{\otimes n}(I_\mu) = \mathbb{C}[\mathbf{x},\mathbf{y}]/(I_\mu(x_1,y_1) + \cdots + I_\mu(x_n,y_n)),$$

the image of $\Delta_L(\mathbf{x},\mathbf{y})$ vanishes for all $L \neq \mu$, while the image of Δ_μ spans the one-dimensional space $\wedge^n B(I_\mu)$. Hence the composite map

$$\mathbb{C}[\mathbf{x},\mathbf{y}] \to B^{\otimes n}(I_\mu) \to \wedge^n B(I_\mu) \cong \mathbb{C}$$

is in effect given by $\Theta^\epsilon(-)|_{\Delta_\mu}$. This yields:

COROLLARY 3.3. *The ideal J_μ is the kernel of the composite map*

$$\mathbb{C}[\mathbf{x},\mathbf{y}] \to B^{\otimes n}(I_\mu) \underset{\phi(I_\mu)}{\to} B^{\otimes n}(I_\mu)^* \otimes \wedge^n B(I_\mu).$$

It might appear that Proposition 3.1 and Corollary 3.3 immediately imply that J_μ is (modulo $I_\mu(x_1,y_1) + \cdots + I_\mu(x_n,y_n)$) the fiber of $\mathcal{J}$ and that $\mathbb{C}[\mathbf{x},\mathbf{y}]/J_\mu$ is the fiber of P at the point $I_\mu \in H_n$. In reality, the matter is not quite so simple. For one thing, we don't know yet that P and $\mathcal{J}$ are locally free, *i.e.*, that they are the sheaves of sections of vector bundles. Absent this, their "fibers" at I_μ are not necessarily what intuition would suggest. Even if we assume that they are locally free, if we carefully examine the passage from sheaves of sections to fibers, the only conclusion we can draw immediately is that J_μ *contains* the fiber $\mathcal{J}(I_\mu)$. Nevertheless, the following facts do hold:

- P (and hence also $\mathcal{J}$) is locally free if and only if X_n is Cohen-Macaulay. This follows because H_n is non-singular and $\rho\colon X_n \to H_n$ is finite; in this situation, the local freeness of $P = \rho_* \mathcal{O}_{X_n}$ is essentially the definition of the Cohen-Macaulay property.
- Generically, P is locally free of rank $n!$, because the map ρ is generically $n!$-to-one. When $I = I(S)$, there are $n!$ choices for the ordering $P_1, \ldots, P_n$ of the distinct points of S.
- It follows that every fiber of the linear map ϕ has rank at most $n!$, and that coker ϕ, P and $\mathcal{J}$ are all locally free in a neighborhood of $I \in H_n$ if and only if $\phi(I)$ has rank $n!$.
- The ring $\mathbb{C}[\mathbf{x},\mathbf{y}]/J_\mu$ is Gorenstein. More generally, a graded Artin algebra $\mathbb{C}[\mathbf{x}]/J$ is Gorenstein if and only if J is the annihilating ideal of a Macaulay inverse system generated by one element [6]. If you like, you may take this as the definition of Gorenstein for a graded Artin algebra; then a general graded algebra is Gorenstein if it is Cohen-Macaulay and its quotient by any (equivalently every) homogeneous parameter ideal is Gorenstein Artin. By the way, Artin means Krull dimension zero, or finite dimension as a vector space.

Using these facts, one can deduce from Proposition 3.1 and Corollary 3.3 the following theorem.

THEOREM 3.4. *Let* $Q_\mu = (I_\mu, \underline{0}, \dots, \underline{0})$ *be the unique point of* X_n *lying over* $I_\mu \in H_n$. *The following are equivalent:*

(1) X_n *is locally Cohen-Macaulay and Gorenstein at* Q_μ*;*
(2) the $n!$ *theorem (Theorem 1.1) holds for the partition* μ.

When these conditions hold, moreover, J_μ *is the ideal of the scheme-theoretic fiber* $\rho^{-1}(I_\mu) \subseteq X_\mu$*, that is, we have* $\mathbb{C}[\mathbf{x}, \mathbf{y}]/J_\mu \cong P(I_\mu)$.

Proof. First assume (1) holds. Then P and $\mathcal{J}$ are locally free, and J_μ contains $\mathcal{J}(I_\mu)$. Moreover, $P(I_\mu) = B^{\otimes n}(I_\mu)/\mathcal{J}(I_\mu)$ is a Gorenstein Artin ring on which S_n acts by the regular representation. Its socle, or unique minimal ideal, is one-dimensional, so S_n must necessarily act on it by the sign representation. If J_μ is not equal to $\mathcal{J}(I_\mu)$, then J_μ must contain the socle, and hence $\mathbb{C}[\mathbf{x}, \mathbf{y}]/J_\mu$ does not contain a copy of the sign representation of S_n. In particular, we must have $\Delta_\mu \in J_\mu$, which is absurd.

Conversely, assume (2) holds. Then $\phi(I_\mu)$ has rank $n!$, so X_n is locally Cohen-Macaulay at I_μ. Furthermore, $\mathcal{J}(I_\mu) = J_\mu$, so $P(I_\mu) \cong \mathbb{C}[\mathbf{x}, \mathbf{y}]/J_\mu$ is Gorenstein, which implies X_n is Gorenstein at I_μ.

We have now tied together two of our seemingly unrelated theorems from Lecture 1, namely, the $n!$ theorem, and the geometric theorem that X_n is Gorenstein. In fact, we have shown that the latter theorem implies the $n!$ theorem.

3.4. THE NESTED HILBERT SCHEME

For ideals $I_{n-1} \in H_{n-1}$, $I_n \in H_n$, the condition that $I_n \subseteq I_{n-1}$ is a closed condition on $H_{n-1} \times H_n$, *i.e.*, it is defined by polynomial equations in local coordinates (much like those for the corresponding condition on a product of Grassmann varieties). Hence we can define the *nested Hilbert scheme*

$$H_{n-1,n} \subseteq \{(I_{n-1}, I_n) : I_n \subseteq I_{n-1}\} \subseteq H_{n-1} \times H_n$$

as a closed subscheme. Geometrically, it parametrizes pairs of finite subschemes $S_{n-1} \subseteq S_n \subseteq \mathbb{C}^2$, of lengths $n-1$ and n. There is a remarkable extension of Fogarty's theorem to the nested case.

THEOREM 3.5. (Tikhomirov, unpublished; see [4] for proof). *The nested Hilbert scheme* $H_{n-1,n}$ *is non-singular and irreducible, of dimension* $2n$.

One can define more general nested Hilbert schemes parametrizing flags of subschemes of various lengths. Cheah [4] has shown that of these *only* H_n and $H_{n-1,n}$ are non-singular, apart from some small exceptions like $H_{1,2,3}$.

We will prove Theorem 1.7, the Gorenstein property of X_n, by induction on n, using the nested Hilbert scheme to provide the geometric link between X_{n-1} and X_n. For this we need to introduce the nested isospectral Hilbert scheme $X_{n-1,n}$. But first we must discuss $H_{n-1,n}$ a little more. Given a pair $I_n \subseteq I_{n-1}$ in $H_{n-1,n}$, there is a unique *distinguished point* $P \in V(I_n)$ such that the length of $(R/I_n)_P$ is exceeds the length of $(R/I_{n-1})_P$ by 1. In terms of the Chow morphisms, if $\sigma_{n-1}(I_{n-1}) = [\![P_1, \ldots, P_{n-1}]\!]$, then $\sigma_n(I_n) = [\![P_1, \ldots, P_{n-1}, P]\!]$, where P is the distinguished point. Another way to look at this is that the R-module I_{n-1}/I_n has length 1, hence is isomorphic to $R/m = \mathbb{C}[x,y]/(x-x_n, y-y_n)$ for some maximal ideal m. Then (x_n, y_n) are the coordinates of P. We get a Chow morphism

$$\sigma_{n-1,n}\colon H_{n-1,n} \to S^{n-1}\mathbb{C}^2 \times \mathbb{C}^2$$

defined by

$$\sigma_{n-1,n}(I_{n-1}, I_n) = (\sigma_{n-1}(I_{n-1}), P).$$

Both maps $H_{n-1,n} \to S^{n-1}\mathbb{C}^2$ and $H_{n-1,n} \to S^n\mathbb{C}^2$ induced by the Chow morphisms σ_{n-1} and σ_n composed with the projections on H_{n-1} and H_n factor through $\sigma_{n-1,n}$.

DEFINITION 3.6. *The nested isospectral Hilbert scheme $X_{n-1,n}$ is the reduced fiber product*

$$\begin{array}{ccc} X_{n-1,n} & \longrightarrow & \mathbb{C}^{2n} \\ \big\downarrow & & \big\downarrow \\ H_{n-1,n} & \longrightarrow & S^{n-1}\mathbb{C}^2 \times \mathbb{C}^2. \end{array} \tag{12}$$

In other words, the points of $X_{n-1,n}$ are tuples $(I_{n-1}, I_n, P_1, \ldots, P_n)$, where $I_n \subseteq I_{n-1}$, $\sigma_n(I_n) = [\![P_1, \ldots, P_n]\!]$, and P_n is the distinguished point. Note that the reduced fiber product Y in the diagram

$$\begin{array}{ccc} Y & \longrightarrow & H_{n-1,n} \\ \big\downarrow & & \big\downarrow \\ X_{n-1} & \longrightarrow & H_{n-1} \end{array} \tag{13}$$

can be identified with the set of tuples $(I_{n-1}, I_n, P_1, \ldots, P_{n-1})$ such that $I_n \subseteq I_{n-1}$ and $\sigma_{n-1}(I_{n-1}) = [\![P_1, \ldots, P_{n-1}]\!]$. Then $X_{n-1,n}$ is the graph of the morphism $Y \to \mathbb{C}^2$ sending $(I_{n-1}, I_n, P_1, \ldots, P_{n-1})$ to its distinguished point $P = P_n$. Thus the above fiber product Y provides an alternative description of $X_{n-1,n}$.

We can now indicate roughly how the proof of Theorem 1.7 goes. We consider the larger diagram

$$\begin{tikzcd}
Y = X_{n-1,n} \arrow[rr, "g"] \arrow[dr] \arrow[dd] & & X_n \arrow[d] \\
& H_{n-1,n} \arrow[r] \arrow[d] & H_n \\
X_{n-1} \arrow[r] & H_{n-1} &
\end{tikzcd} \tag{14}$$

We are to prove that X_n is Gorenstein, and we will assume the result for X_{n-1} by induction. Then the bottom arrow in (14) is flat with Gorenstein fibers, since it is finite and H_{n-1} is non-singular. Hence the same holds for the diagonal arrow if we take Y to be the *scheme-theoretic* fiber product. Since $H_{n-1,n}$ is non-singular, the scheme-theoretic fiber product is again Gorenstein. It is generically reduced, hence reduced, hence equal to $X_{n-1,n}$. This shows that $X_{n-1,n}$ is Gorenstein. The morphism $g : X_{n-1,n} \to X_n$ given by the top arrow in (14) is projective and birational, allowing $X_{n-1,n}$ to play the role of a partial desingularization of X_n. To complete the proof we will use duality for the morphism g to deduce that X_n is Gorenstein, as desired. In order to do this we will need the identity

$$Rg_*\mathcal{O}_{X_{n-1,n}} = \mathcal{O}_{X_n}$$

for the *derived functor* Rg_* of the sheaf pushforward functor g_*. At the end of the day, to establish this identity we will need know something more about X_n than what follows merely from the inductive set-up. What that something is and how to prove it will be taken up in the next lecture.

4. The main theorem and the polygraph theorem

4.1. MAIN THEOREM

We start by adding some details to the outline of the proof of Theorem 1.7 given at the end of the last lecture. We want to prove that X_n is Gorenstein by induction on n, but we need to refine the induction hypothesis a bit.

The Proj construction gives rise to a line bundle denoted $\mathcal{O}(1)$ on any $X = \operatorname{Proj} S$, such that the elements of S_1 represent global sections of $\mathcal{O}(1)$. The k-th tensor power of $\mathcal{O}(1)$ is denoted $\mathcal{O}(k)$; if k is negative, this means the $(-k)$-th tensor power of the dual $\mathcal{O}(-1) = \mathcal{O}(1)^*$. In the case of H_n

and X_n, it follows from the proof of Theorem 2.1 that $\mathcal{O}(1) = \wedge^n B$. Every Gorenstein scheme X possesses a *canonical line bundle* (more on this below) which in the case of a non-singular scheme is the usual canonical line bundle, *i.e.*, the sheaf of exterior differential d-forms, where $d = \dim X$. What we will prove by induction is that

X_n is Gorenstein, with canonical line bundle equal to $\mathcal{O}(-1)$.

As explained in the last lecture, we can deduce from the induction hypothesis for $n-1$ that $X_{n-1,n}$, is Gorenstein. In fact we can do a little better. The *relative* canonical line bundles for the bottom arrow and the diagonal arrow in (14) are the same. On $H_{n-1,n}$ and $X_{n-1,n}$, we write $\mathcal{O}(k,l)$ for the tensor product of the pullback of $\mathcal{O}(k)$ from H_{n-1} with the pullback of $\mathcal{O}(l)$ from H_n. By an explicit calculation in [17], the canonical sheaf of H_{n-1} is $\mathcal{O}$ (this was well-known) and that of $H_{n-1,n}$ is $\mathcal{O}(1,-1)$ (this is new). By the induction hypothesis and the calculation for H_{n-1}, the relative canonical sheaf for the bottom and diagonal arrows in (14) is $\mathcal{O}(-1,0)$. Hence by the calculation for $H_{n-1,n}$ the canonical sheaf on $X_{n-1,n}$ is $\mathcal{O}(0,-1)$, that is, the pullback $g^*\mathcal{O}_{X_n}(-1)$ of $\mathcal{O}(-1)$ from X_n.

The key to the proof is to show that for the projection $g\colon X_{n-1,n} \to X_n$ we have

$$Rg_*\mathcal{O}_{X_{n-1,n}} = \mathcal{O}_{X_n}.$$

Here Rg_* is the *derived functor* of the pushforward g_*. For present purposes, everything you need to know about derived functors and derived categories can be summed up as follows.

- On every scheme X of finite type over a field there is a complex of sheaves ω_X, which is a *dualizing object* in the derived category. This means that the derived functor $\mathcal{D}_X = R\mathcal{H}\mathrm{om}(-,\omega_X)$ satisfies $\mathcal{D}_X \circ \mathcal{D}_X = \mathrm{id}$. Although a dualizing object is not unique, there is a canonical preferred choice (up to isomorphism in the derived category).
- If X is Cohen-Macaulay and d-dimensional, then ω_X reduces to a single sheaf concentrated in degree $-d$. If X is also Gorenstein, that sheaf is the canonical line bundle on X. The converse statements also hold.
- *Duality theorem*: if $g\colon Y \to X$ is a proper (*e.g.*, projective) morphism, then the derived functor Rg_* commutes with duality, *i.e.*, $\mathcal{D}_X \circ Rg_* = Rg_* \circ \mathcal{D}_Y$.
- *Projection formula*: if L is a locally free sheaf on X, then $Rg_*(A \otimes g^*L) = (Rg_*A) \otimes L$.

Now let's grant for a moment that $Rg_*\mathcal{O}_{X_{n-1,n}} = \mathcal{O}_{X_n}$. We have established that $X_{n-1,n}$ is Gorenstein with canonical sheaf $g^*\mathcal{O}_{X_n}(-1)$, and hence its dualizing complex is this sheaf, concentrated in degree $-2n$. The morphism

g is projective. On the one hand, using the projection formula with $A = \mathcal{O}_{X_{n-1,n}}$ and $L = \mathcal{O}_{X_n}(-1)$, we deduce

$$Rg_* g^* \mathcal{O}_{X_n}(-1) = \mathcal{O}_{X_n}(-1).$$

On the other hand, using the duality theorem, we deduce

$$Rg_* \omega_{X_{n-1,n}} = \omega_{X_n}.$$

Together, these two identites show that ω_{X_n} is the sheaf $\mathcal{O}_{X_n}(-1)$ concentrated in degree $-2n$. Hence X_n is Gorenstein with canonical sheaf $\mathcal{O}(-1)$, which is what we wanted to prove.

How do we show that $Rg_* \mathcal{O}_{X_{n-1,n}} = \mathcal{O}_{X_n}$? First of all, this is a local question on the target X_n of g. On the locus where $P_1, \dots, P_n$ are not all equal, X_n is locally isomorphic to a product $X_k \times X_l$ for some $k, l < n$, and g is locally isomorphic to $g'\colon X_k \times X_{l-1,l} \to X_k \times X_l$. This is intuitively easy to understand and not hard to prove, so I refer you to [17] for the details. We can assume as part of the induction that $Rg'_* \mathcal{O} = \mathcal{O}$, so we only have to worry about the situation in the neighborhood of points of X_n where $P_1, \dots, P_n$ all coincide. In particular, we only have to worry about what happens where $y_1 - y_2$, $y_2 - y_3$, $\dots$, $y_{n-1} - y_n$ all vanish.

Now there is a standard technique for dealing with this situation using *local cohomology.* Again I refer you to [17] for technical details, but the upshot is that to extend the identity $Rg_* \mathcal{O}_{X_{n-1,n}} = \mathcal{O}_{X_n}$ from the complement of $V(y_1 - y_2, \dots, y_{n-1} - y_n)$ to the whole of X_n and $X_{n-1,n}$ we only need to establish three facts:

(a) The sequence $y_1 - y_2, \dots, y_{n-1} - y_n$ is a regular sequence in the local ring of $\mathcal{O}_{X_n}$ at every point $x \in V(y_1 - y_2, \dots, y_{n-1} - y_n)$,
(b) The same holds in $\mathcal{O}_{X_{n-1,n}}$, and
(c) The fibers of g have dimension less than $n - 2$.

To complete your confusion, we define another morphism

$$\alpha\colon H_{n-1,n} \to F \tag{15}$$

sending (I_{n-1}, I_n) to (I_n, P). Since the distinguished point can be any point of I_n (for some choice of I_{n-1}), the image of α is indeed F. Generically, when $I_n = I(S)$, we must have $I_{n-1} = I(S \setminus \{P\})$, so the fiber of α over $(I(S), P)$ has just one point. Thus α is birational (it's also projective). The fibers of g are actually fibers of α, as you will see if you understand the remarks following Definition 3.6. The fiber of α over a point (I, P) is the projective space $\mathbb{P}(A)$ where A is the socle of the local ring $(R/I)_P$, that is, the annihilator of the maximal ideal. Its dimension must be maximized at a monomial ideal. Hence the maximum fiber dimension is one less than the maximum number of corners of the diagram of a partition μ of n. For $n \geq 4$, the number of corners of μ is always less than $n-1$, so fact (c) holds.

For fact (b), we already know that $X_{n-1,n}$ is Cohen-Macaulay, so we only have to show that the locus $V(y_1-y_2, \ldots, y_{n-1}-y_n)$ has codimension $n-1$ in $X_{n-1,n}$. Then it is a local complete intersection, and, by the general theory of Cohen-Macaulay rings, the generators of its ideal form a regular sequence. The required dimension estimate follows from the cell decomposition of $H_{n-1,n}$ in [4].

The heart of the matter is fact (a). It is true that $V(y_1-y_2, \ldots, y_{n-1}-y_n)$ has codimension $n-1$ in X_n, but this information is useless since we don't know in advance that X_n is Cohen-Macaulay. What we need is the following lemma.

LEMMA 4.1. *The map* $X_n \to \mathbb{C}^n = \operatorname{Spec} \mathbb{C}[\mathbf{y}]$, *given by the composite of the blowup map* $X_n \to \mathbb{C}^{2n}$ *with the projection* $\mathbb{C}^{2n} \to \mathbb{C}^n$ *on the* $\mathbf{y}$ *coordinates, is flat.*

Granting this lemma, the proof is otherwise complete, except that we only have fact (c) for $n \geq 4$. For $n = 1, 2, 3$, we need to verify the induction hypothesis directly. The action of $(\mathbb{C}^*)^2$ on $\mathbb{C}^2$ as the group of 2×2 diagonal matrices induces an action of $(\mathbb{C}^*)^2$ on X_n in which every point has some Q_μ in the closure of its orbit. Using this fact, we can reduce the problem of showing that X_n is Gorenstein to the local problem at the distinguished points Q_μ. Then, using Theorem 3.4, we need only verify the $n!$ theorem for each Δ_μ with $|\mu| \leq 3$. This is an easy computation which I invite you to carry out by hand. The only case you really have to check is $\mu = (2, 1)$, since the others are classical Vandermonde determinants in one set of variables.

The multiplication and alternation pairing $B^{\otimes n} \otimes B^{\otimes n} \to \wedge^n B = \mathcal{O}(1)$ in Proposition 3.1 induces a pairing $P \otimes P \to \mathcal{O}(1)$. Once we know X_n is Gorenstein, P becomes a vector bundle and the induced pairing is perfect on every fiber. By the duality theorem for the finite morphism $\rho\colon X_n \to H_n$ this implies that $\omega_{X_n} = \mathcal{O}(-1)$. So we have verified the full induction hypothesis, as we stated it, for $n = 1, 2, 3$.

Technically, the identity $Rg_*\mathcal{O}_{X_{n-1,n}} = \mathcal{O}_{X_n}$ is part of the induction hypothesis too. For $n = 1, 2$, the morphism g is an isomorphism, and this identity is trivial. For $n = 3$, g is everywhere locally an isomorphism, except over the locus Z consisting of non-curvilinear triple points. But in this case, we know that both $X_{2,3}$ and X_3 are Cohen-Macaulay, and we can find a sequence z_1, z_2, z_3 of global functions, regular on both $X_{2,3}$ and X_3 and vanishing on Z. Using this we can verify that $Rg_*\mathcal{O}_{X_{2,3}} = \mathcal{O}_{X_3}$ by a standard local cohomology argument, like the one alluded to earlier in our discussion of the induction step.

4.2. THE POLYGRAPH THEOREM ENTERS THE SCENE

Recall from Lecture 1 the definition of the polygraph $Z(n,l) \subseteq \mathbb{C}^{2n+2l}$ and the theorem (whose proof we will discuss later on) that its coordinate ring is a free $\mathbb{C}[\mathbf{y}]$-module, where $\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b}$ are the coordinates on $\mathbb{C}^{2n+2l}$. We will now show how Lemma 4.1 follows from the polygraph theorem and the construction of X_n as a blowup. Recall specifically that

$$X_n = \operatorname{Proj} \mathbb{C}[\mathbf{x}, \mathbf{y}][tJ] = \operatorname{Proj} \bigoplus_d J^d,$$

where $J = \mathbb{C}[\mathbf{x}, \mathbf{y}]A$ is the ideal generated by the alternating polynomials. To establish Lemma 4.1 it suffices to prove the following proposition.

PROPOSITION 4.2. *The ideal J^d is a free $\mathbb{C}[\mathbf{y}]$-module for all d.*

Proof. Set $l = dn$, and let $Z(n,l) \subseteq \mathbb{C}^{2n+2l}$ be the polygraph. Recall that this is the union of the subspaces

$$W_f = V(a_i - x_{f(i)}, b_i - y_{f(i)} : 1 \leq i \leq l),$$

for all $f : \{1, \ldots, l\} \to \{1, \ldots, n\}$. Let $G = S_n^d$ be the Cartesian product of d copies of the symmetric group S_n, acting on $\mathbb{C}^{2n+2l}$ by permuting the factors in $\mathbb{C}^{2l} = (\mathbb{C}^2)^{dn}$ in d consecutive blocks of length n. In other words each $w \in G$ fixes the coordinates $\mathbf{x}, \mathbf{y}$ on $\mathbb{C}^{2n}$, and for each $k = 0, \ldots, d-1$ permutes the coordinate pairs a_{kn+1}, b_{kn+1} through a_{kn+n}, b_{kn+n} among themselves.

Let $R(n,l) = \mathbb{C}[\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b}]/I(n,l)$ be the coordinate ring of $Z(n,l)$. By Theorem 1.8, $R(n,l)$ is a free $\mathbb{C}[\mathbf{y}]$-module. By the symmetry of its definition, $I(n,l)$ is a G-invariant ideal, so G acts on $R(n,l)$. We claim that J^d is isomorphic as a $\mathbb{C}[\mathbf{x}, \mathbf{y}]$-module to the space $R(n,l)^\epsilon$ of G-alternating elements of $R(n,l)$. Each x-degree homogeneous component of $R(n,l)$ is a finitely generated y-graded free $\mathbb{C}[\mathbf{y}]$-module. Since $R(n,l)^\epsilon$ is a graded direct summand of $R(n,l)$, it is a free $\mathbb{C}[\mathbf{y}]$-module, so the claim proves the Proposition.

Define a function f_0 by $f_0(kn+i) = i$ for all $0 \leq k < d$, $1 \leq i \leq n$. Restriction of regular functions from $Z(n,l)$ to its component subspace W_{f_0} is given by the $\mathbb{C}[\mathbf{x}, \mathbf{y}]$-algebra homomorphism $\psi : R(n,l) \to \mathbb{C}[\mathbf{x}, \mathbf{y}]$ mapping a_{kn+i}, b_{kn+i} to x_i, y_i. Observe that ψ maps $R(n,l)^\epsilon$ surjectively onto $\mathbb{C}[\mathbf{x}, \mathbf{y}]A^d = J^d$. In fact, $R(n,l)^\epsilon$ is the $\mathbb{C}[\mathbf{x}, \mathbf{y}]$-submodule of $R(n,l)$ generated by products of the form

$$\Delta_{M_1}(a_1, b_1, \ldots, a_n, b_n) \cdots \Delta_{M_d}(a_{(d-1)n+1}, b_{(d-1)n+1}, \ldots, a_{dn}, b_{dn}).$$

Let p be an arbitrary element of $R(n,l)^{\epsilon}$. Since p is G-alternating, p vanishes on W_f if $f(kn+i) = f(kn+j)$ for some $0 \leq k < d$ and some $1 \leq i < j \leq n$. Thus the regular function defined by p on $Z(n,l)$ is determined by its restriction to those components W_f such that for each k, the sequence $f(kn+1), \ldots, f(kn+n)$ is a permutation of $\{1, \ldots, n\}$. Moreover, for every such f there is an element $w \in G$ carrying W_f onto W_{f_0}. Hence p is determined by its restriction to W_{f_0}. This shows that p vanishes on $Z(n,l)$ if $\psi(p) = 0$, that is, the kernel of $\psi\colon R(n,l)^{\epsilon} \to J^d$ is zero. Hence ψ is injective as well surjective, so it's an isomorphism. □

At this point, everything is left hanging solely on the proof of the polygraph theorem. Before we turn to the discussion of that proof, I want to make a few more remarks on the geometric topics.

The first remark is that the polygraph has a *direct* geometric relevance to the Hilbert scheme picture, which is somewhat obscured by its roundabout entrance via Lemma 4.1 and Proposition 4.2. To see this, consider the universal family $F/X_n = F \times_{H_n} X_n$ pulled back to X_n, and its l-th fiber power

$$F^l/X_n = \{(I, P_1, \ldots, P_n, Q_1, \ldots, Q_l) : Q_i \in V(I) \text{ for all } i\}.$$

Since $\sigma(I) = [\![P_1, \ldots, P_n]\!]$ the condition $Q_i \in V(I)$ is equivalent to $Q_i \in \{P_1, \ldots, P_n\}$. It follows immediately that the projection on $\mathbb{C}^{2n+2l}$ of $F^l/X_n \subseteq H_n \times \mathbb{C}^{2n+2l}$ is nothing other than the polygraph $Z(n,l)$. In particular, every regular function on $Z(n,l)$ can be composed with this projection to yield a global regular function on F^l/X_n. The latter is the same thing as a global section of $B^{\otimes l}$ over X_n (abusing notation, we write B for what is really $\rho^* B$). So we have a canonical, geometrically defined ring homomorphism

$$R(n,l) \to H^0(X_n, B^{\otimes l}).$$

Later we will see that this homomorphism is an isomorphism, and what's more, the higher cohomology modules $H^i(X_n, B^{\otimes l})$ vanish for $i > 0$. If we grant in advance that X_n is Cohen-Macaulay, and observe that the locus $V(\mathbf{y})$ is a complete intersection in X_n, these facts in turn *imply* that $R(n,l)$ is a free $\mathbb{C}[\mathbf{y}]$-module. So, conceptually, the right idea is that $R(n,l)$ is free *because* X_n is Cohen-Macaulay and the tensor powers of B enjoy a strong cohomology vanishing theorem on X_n. For now, however, we are obliged to work backwards and prove that $R(n,l)$ is free by getting our hands dirty.

The second remark concerns the ideal J and its powers. Any alternating polynomial must obviously vanish at a point of $\mathbb{C}^{2n}$ where $x_i = x_j$ and $y_i = y_j$, for $i \neq j$. Thus we have obvious inclusions

$$J \subseteq \bigcap_{i<j} (x_i - x_j, y_i - y_j),$$

and consequently

$$J^d \subseteq \bigcap_{i<j} (x_i - x_j, y_i - y_j)^d.$$

It is a consequence of Proposition 4.2 that equality holds here:

PROPOSITION 4.3. *For all n and d, we have*

$$J^d = \bigcap_{i<j} (x_i - x_j, y_i - y_j)^d. \tag{16}$$

Proof. We first consider the situation locally at a point where the y_i are not all equal. With a little work, it can be shown that (16) reduces at such a point to the same equation for smaller values of n. Thus we can assume by induction that (16) holds locally on $U = \mathbb{C}^{2n} \setminus V(y_1 - y_2, \ldots, y_{n-1} - y_n)$. Since J^d is a free $\mathbb{C}[\mathbf{y}]$-module, the ring $\mathbb{C}[\mathbf{x}, \mathbf{y}]/J^d$ has depth at least $n-1$ as a $\mathbb{C}[\mathbf{y}]$-module. This implies that $V(J^d)$ is scheme-theoretically equal to the closure of its intersection with any open set of the form $\mathbb{C}^{2n} \setminus V(I)$, where I is generated by at least two independent linear forms in $\mathbf{y}$. In particular, for $n \geq 3$, $V(J^d)$ is the closure of its intersection with U, establishing (16). For $n = 1, 2$, the result is trivial. □

Proposition 4.3 and Theorem 2.1 justify the assertion that H_n and X_n are the blowups of $S^n\mathbb{C}^2$ and $\mathbb{C}^{2n}$, respectively, along the reduced union of pairwise diagonals. See the remarks following Theorem 2.1. Proposition 4.3 also implies that the ideals J^d are integrally closed, and hence the Reese algebra $T = \mathbb{C}[\mathbf{x}, \mathbf{y}][tJ]$ is an integrally closed domain. By definition, this shows that X_n is *arithmetically normal* in its realization as $\operatorname{Proj} T$. In particular, X_n is normal, *i.e.*, its local ring at every point is integrally closed.

We conjecture that Theorem 1.8 holds for the polygraph $Z(n, l) \subseteq \mathbb{C}^{kn+kl}$ over $E = \mathbb{C}^k$ for any k, with freeness over the polynomial ring $\mathbb{C}[\mathbf{z}]$ in any one set of n coordinates on $\mathbb{C}^{kn}$. If true, this would imply that Proposition 4.2 and Proposition 4.3 hold in more than two sets of variables $\mathbf{x}, \mathbf{y}, \ldots, \mathbf{z}$, with essentially the same proof. At present we cannot prove any of these results for $k > 2$, not even the "intuitively obvious" case $d = 1$ of Proposition 4.3. Indeed we cannot prove these results for $k = 2$ without recourse to the polygraph theorem. For $k = 1$, by contrast, they are easy and well-known: the ideal J is generated by the Vandermonde determinant $\Delta(\mathbf{x})$, and the intersection of ideals in (16) is equal to the their product.

4.3. THE POLYGRAPH BASIS

Rather than go over the proof of the polygraph theorem as presented in [17], what I would like to do is reveal the secret behind it. The secret is an algorithm for producing the elements of a free module basis of $R(n,l)$, in terms of combinatorial indexing data. The indexing data are pairs

$$e = (e_1, \ldots, e_n) \in \mathbb{N}^n, \quad f\colon [l] \to [n].$$

Here and below we use the abbreviation $[k] = \{1, \ldots, k\}$. Recall that $R(n,l)$ is a quotient of a polynomial ring

$$R(n,l) = \mathbb{C}[\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b}]/I(n,l),$$

so our basis elements will be represented by polynomials

$$p[e,f] \in \mathbb{C}[\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b}].$$

As $R(n,l)$ is doubly graded by degree in the $\mathbf{x}, \mathbf{a}$ variables ("x-degree") and the $\mathbf{y}, \mathbf{b}$ variables ("y-degree"), our elements $p[e,f]$ will be doubly homogeneous. They will also satisfy certain important vanishing conditions. To describe the conditions, we introduce loci $Y(m,r,k) \subseteq Z(n,l)$ defined by

$$Y(m,r,k) = \bigcup V(x_T) \cap W_f : \quad |T \cap [r] \setminus f([k])| \geq m.$$

Here T is a subset of $[n]$ and $V(x_T)$ is shorthand for $V(x_j : j \in T)$. Recall that W_f is the component of the polygraph $Z(n,l)$ on which $a_i = x_{f(i)}$, $b_i = y_{f(i)}$ for all $i \in [l]$. Thus, informally, $Y(m,r,k)$ is the locus on which at least m of the first r coordinates x_j vanish, coordinates in $f([k])$ excluded. The ideal of $Y(m,r,k)$ is called $I(m,r,k)$. The basis elements $p[e,f]$ will satisfy

- $\deg_x p[e,f] = |e| = e_1 + \cdots + e_n$,
- $\deg_y p[e,f] = i(e,f)$, defined below,
- $p[e,f] \in I(m,r,k)$ iff $|(\operatorname{supp} e)^c \cap [r] \setminus f([k])| < m$, and these elements span $I(m,r,k)$ as a $\mathbb{C}[\mathbf{y}]$-module.

Here $\operatorname{supp} e$ is the set of indices j with $e_j \neq 0$, and $(\operatorname{supp} e)^c$ is its complement. The last condition implies that the ideals $I(m,r,k)$ generate a distributive sublattice of the lattice of ideals in $R(n,l)$, all of whose members are spanned by subsets of the basis $\{p[e,f]\}$. It follows that the coordinate ring of any intersection of unions of the loci $Y(m,r,k)$ is again a free $\mathbb{C}[\mathbf{y}]$-module (and also that every such intersection is scheme-theoretically reduced).

Let's check that at least in a certain weak sense, the numbers work out. Any free $\mathbb{C}[\mathbf{y}]$-module basis of $R(n,l)$ must also be a $\mathbb{C}(\mathbf{y})$-vector space basis of $\mathbb{C}(\mathbf{y}) \otimes R(n,l)$. In $\mathbb{C}(\mathbf{y}) \otimes R(n,l)$ we can form for each f the rational function

$$1_f = \prod_i \prod_{j \neq f(i)} \frac{b_i - y_j}{y_{f(i)} - y_j},$$

which is identically 1 on W_f and vanishes on all other components W_g (it is undefined where W_f and W_g intersect). The set of all elements

$$x^e 1_f$$

is a $\mathbb{C}(\mathbf{y})$-basis of $\mathbb{C}(\mathbf{y}) \otimes R(n,l)$. In this basis, it is not hard to see that $\mathbb{C}(\mathbf{y}) \otimes I(m,r,k)$ is spanned by the elements

$$x^e 1_f : \quad |(\operatorname{supp} e)^c \cap [r] \setminus f([k])| < m.$$

The reason is that if $|T \cap [r] \setminus f([k])| \geq m$ and e, f are as above, then $(\operatorname{supp} e)^c$ cannot contain T, that is, $(\operatorname{supp} e) \cap T$ is non-empty, so x^e must vanish on $V(x_T)$. This shows that in each x-degree, we have specified the right number of elements $p[e,f]$ to belong in $I(m,r,k)$.

Before describing the algorithm, I want to introduce a diagram formalism for the indexing data e, f, and define $i(e,f)$. We associate to the pair e, f an $n \times l$ *grid* with numerical entries. The numbers $e_1, \ldots, e_n$ are placed along the left edge of the grid. In column i, we place a number in row $f(i)$, choosing these numbers to increase consecutively in each row j, starting with e_j at the left. Here's an example: let $n = 4$, $l = 6$,

$$e = (0,5,1,3), \quad f = (4,4,1,2,1,3).$$

The associated grid is

$$\begin{array}{r|c|c|c|c|c|c|}
\cline{2-7}
3 & 4 & 5 & & & & \\
\cline{2-7}
1 & & & & & & 2 \\
\cline{2-7}
5 & & & & 6 & & \\
\cline{2-7}
0 & & & 1 & & 2 & \\
\cline{2-7}
\end{array}.$$

The "inversion number" $i(e,f)$ is now defined to be the number of pairs of entries in distinct rows and columns (including the left edge as a column) satisfying either of two conditions:

- the right entry is in a higher row than the left entry, and the two entries have equal value, or
- the right entry is in a lower row and is equal to the left entry plus one.

In the example there are $i(e,f) = 4$ such pairs. They are the two 2's, the two 5's, the 1 on the left edge with the lower 2, and the upper 5 with the 6.

If we define the generating function

$$A^{n,l}(q,t) = \frac{1}{(1-q)^n} \sum_{e,f} t^{|e|} q^{i(e,f)},$$

it follows that the Hilbert series of $R(n,l)$ as a doubly graded algebra is given by

$$\sum_{r,s} t^r q^s \dim R(n,l)_{r,s} = A^{n,l}(q,t).$$

In particular we must have $A^{n,l}(q,t) = A^{n,l}(t,q)$, a fact which is rather amazing when considered in light of the combinatorial definition. Later we will see that there is an explicit formula for $A^{n,l}(q,t)$ in terms of operators on Macdonald polynomials.

The quantity $i(e,f)$ can be computed by a combinatorial recurrence. The key to this is to observe that $i(e,f)$ remains invariant if we "cycle" the grid by moving the bottom row to the top, and subtracting one from each entry of that row. Doing this only gives a new legal grid if $e_1 > 0$, but there are modifications that work when $e_1 = 0$. Altogether there are three cases. Once and for all we fix the cyclic permutation

$$\theta = (1\,2\,\ldots\,n).$$

In terms of θ the result of cycling the grid for e, f is the grid for $e' = \theta^{-1}e - (0,\ldots,0,1)$, $f' = \theta^{-1}f$.

Case I: $e_1 = 0$ and $1 \notin f([l])$, *i.e.* the only entry in the bottow row is a zero at the left edge. The bottom row contributes nothing to $i(e,f)$, so $i(e,f) = i(e',f')$, where $e' = (\theta^{-1}e)|_{[n-1]}$, $f' = \theta^{-1}f$. Note that $e' \in \mathbb{N}^{n-1}$ and f' has image in $[n-1]$. The grid of e', f' is gotten by deleting the bottom row from the grid of e, f.

Case II: $e_1 = 0$ and $t = \min f^{-1}(\{1\})$, *i.e.* there is a 1 in the bottom row of the grid, in column t. Then

$$i(e,f) = i(e',f') + m,$$

where $e' = \theta^{-1}e$, $f' = \theta^{-1}f|_L$, $L = [l]\setminus\{t\}$, and $m = |(\operatorname{supp} e)^c \setminus f^{-1}([t])|$. In other words, m is the number of rows in the grid with a zero at the left edge, and no entry in columns $1,\ldots,t$ (in particular, the bottom row does *not* count). The grid of e', f' is gotten by deleting column t from the grid of e, f and cycling *without* diminishing e_1. Note however that any remaining entries in the bottom row will get diminished as a result of deleting column t.

Case III: $e_1 > 0$. Then $i(e,f) = i(e',f')$, where $e' = \theta^{-1}e - (0,\ldots,0,1)$, $f' = \theta^{-1}f$. The grid of e', f' is the result of cycling the grid of e, f, diminishing e_1 by 1.

Note that the quantity $|e|+n+l$ is reduced by 1 at every step, so the recurrence terminates with the empty grid, defined to have $i(e,f)=0$. The example considered above belongs to Case II, and for this example the grid of e', f' is

$$\begin{array}{c|c|c|c|c|c|} \cline{2-6} 0 & & & & 1 & \\ \cline{2-6} 3 & 4 & 5 & & & \\ \cline{2-6} 1 & & & & & 2 \\ \cline{2-6} 5 & & & 6 & & \\ \cline{2-6} \end{array}.$$

Here $m=0$, since only the original bottom row had $e_j=0$, so the value $i(e,f)$ is equal to that of the new diagram, also 4.

Exercise: Verify that $i(e,f)$ satisfies the recurrence. Cases I and III are trivial, but Case II is not.

Exercise: Let us define the generating function

$$A^{n,l}_{m,r,k}(q,t)=\frac{1}{(1-q)^n}\sum_{e,f} t^{|e|}q^{i(e,f)},$$

summed over pairs e,f with $|(\operatorname{supp} e)^c\cap[r]\setminus f([k])|\geq m$. Since $I(m,r,k)$ is spanned by the elements $p[e,f]$ for e,f not in the range of summation, the Hilbert series of the coordinate ring of $Y(m,r,k)$ is given by

$$\sum_{r,s} t^r q^s \dim(R(n,l)/I(m,r,k))_{r,s}=A^{n,l}_{m,r,k}(q,t).$$

Note that giving the numbers $|(\operatorname{supp} e)^c\cap[r]\setminus f([k])|$ for all r and k is equivalent to giving the sets $S_k(e,f)=(\operatorname{supp} e)\cup f([k])$ for all k. Let us also define generating functions

$$B^{S_0,\ldots,S_l}_n(q,t)=\frac{1}{(1-q)^n}\sum_{e,f} t^{|e|}q^{i(e,f)}$$

summed over pairs e,f with $S_k(e,f)=S_k$ for all $k=0,1,\ldots,l$, so we have

$$A^{n,l}_{m,r,k}(q,t)=\sum_{|[r]\setminus S_k|\geq m} B^{S_0,\ldots,S_l}_n,$$

summed over admissible sequences of sets $\underline{S}$, *i.e.*, sequences $S_0\subseteq\cdots\subseteq S_l\subseteq[n]$ such that $|S_k\setminus S_{k-1}|\leq 1$ for all $k=1,\ldots,l$. Now show that the recurrence for $i(e,f)$ implies that the generating functions $B^{S_0,\ldots,S_l}_n$ are determined by the following recurrence (together with the initial condition $B^{\emptyset,\ldots,\emptyset}_0=\delta_{0,l}$). If $1\notin S_l$, then

$$B^{S_0,\ldots,S_l}_n=\frac{1}{1-q}B^{\theta^{-1}S_0,\ldots,\theta^{-1}S_l}_{n-1}.$$

If $1 \in S_t \setminus S_{t-1}$, then

$$B_n^{S_0,\ldots,S_l} = q^{n-|S_t|} \sum_{S'_0,\ldots,S'_{l-1}} B_n^{S'_0,\ldots,S'_{l-1}},$$

summed over admissible $\underline{S}'$ such that $S'_i = \theta^{-1} S_i$ for $i < t$ and $S'_i \cup \{n\} = \theta^{-1} S_{i+1}$ for $i \geq t$ (note that $n \in \theta^{-1} S_{i+1}$ for these values of i). If $1 \in S_0$, then

$$B_n^{S_0,\ldots,S_l} = t \sum_{S'_0,\ldots,S'_l} B_n^{S'_0,\ldots,S'_l},$$

where $\underline{S}'$ is admissible and $S'_i \cup \{n\} = \theta^{-1} S_i$ for all i.

5. Constructing the basis

5.1. BASIS CONSTRUCTION ALGORITHM

Recall from the last lecture that our plan is to give a basis of $R(n, l)$ consisting of doubly homogeneous polynomials $p[e, f]$ for $e \in \mathbb{N}^n$ and $f\colon [l] \to [n]$ such that

- $\deg_x p[e, f] = |e|$,
- $\deg_y p[e, f] = i(e, f)$,
- $p[e, f] \in I(m, r, k) \Leftrightarrow |(\operatorname{supp} e)^c \cap [r] \setminus f([k])| < m$,

where $i(e, f)$ is the number of inversions in the grid diagram of (e, f), as defined in the last lecture. Also, the elements $p[e, f]$ we have specified to belong to $I(m, r, k)$ will actually span this ideal as a $\mathbb{C}[\mathbf{y}]$-module.

We now present an algorithm to compute $p[e, f]$. In places, the algorithm specifies that we should compute some polynomial as the unique solution of some linear equations over $\mathbb{C}[\mathbf{y}]$. At such points, of course, there is something to prove. First, the solution should exist and be unique in linear algebra over the field of rational functions $\mathbb{C}(\mathbf{y})$, and second, the resulting polynomial should belong to $\mathbb{C}[\mathbf{x}, \mathbf{y}, \mathbf{a}, \mathbf{b}]$ when, *a priori*, it might only belong to $\mathbb{C}(\mathbf{y})[\mathbf{x}, \mathbf{a}, \mathbf{b}]$. These things are true, but I will not attempt to justify them fully here.

You can readily imagine a possible proof of the polygraph theorem along these lines: (1) present the algorithm, (2) prove that it works. The approach actually taken in [17] looks quite different. There I inductively construct bases for the rings $R(n, l)/I(1, 1, t)$ for t going from l to 0, and finally for $R(n, l)$. The algorithm presented here is what you get by unravelling the inductive construction presented in Section 4.10 of [17], and the proof that it works is implicit in the proof of the relevant lemmas there. After presenting the algorithm, I'll say a bit below about what is going on in the proof of its correctness, and why the less algorithmic approach in [17] seems to be the only workable method.

In the algorithm to compute basis elements $p[e,f]$ in $R(n,l)$, we will assume that they have already been computed in $R(n-1,l)$, in $R(n,l-1)$, and in smaller x-degrees in $R(n,l)$. The starting point for the recurrence is the case $l=0$. Then we have $R(n,0)=\mathbb{C}[\mathbf{x},\mathbf{y}]$, and the required basis is

$$p[e,\emptyset]=x^e.$$

This also covers the case $n=0$, since $R(0,l)=0$ if $l>0$.

Given (e,f), we set

$$t=\begin{cases} 0 & \text{if } e_1>0 \\ l+1 & \text{if } e_1=0 \text{ and } 1\notin f([l]) \\ \min f^{-1}(\{1\}) & \text{otherwise.} \end{cases}$$

The value of t is, so to speak, the position of 1 in the bottom row of the grid, with the convention that a missing 1 is "outside" the grid, on the left edge if $e_1>1$ and on the right edge if $e_1=0$. We fix

$$\theta=(1\ 2\ \ldots\ n),$$

as before.

The algorithm proceeds in two phases. In the first phase, we construct a *proto-basis*, whose elements we denote $q[e,f]$. In the second phase, we correct these to get the $p[e,f]$. To compute an individual $p[e,f]$ we would first have to compute all the proto-basis elements $q[e,f]$ of the same bi-degree, so in practice we might as well compute all $p[e,f]$ of a given bi-degree at once. The proto-basis construction has three cases. Comparing them with the three cases in the recurrence for $i(e,f)$, you will see that $i(e,f)$ does correctly give the y-degree of the constructed elements.

Case I: $t=l+1$. In other words, $e=(0,e_2,\ldots,e_n)$ and f maps $[l]$ into $[n]\setminus\{1\}$. Note that $(\theta^{-1}e)|_{[n-1]}\in\mathbb{N}^{n-1}$ determines e and that $\theta^{-1}f$ maps $[l]$ into $[n-1]$. In this case, set

$$q[e,f]=\theta p_{n-1,l}[(\theta^{-1}e)|_{[n-1]},\theta^{-1}f],$$

the right-hand side having already been computed in $R(n-1,l)$.

Case II: $1\le t\le l$. Define

$$Z'(r,t)=V(a_t)\cap\bigcup_{f'}W_{f'}:\quad f'(t)\in[r]\setminus f'([t-1]),\tag{17}$$

a subspace arrangement contained in $R(n,l)$ as a closed subset. Set $L=[l]\setminus\{t\}$, and denote by $R(n,L)$ the ring $R(n,l-1)$ using $\mathbf{a},\mathbf{b}$ variables

$a_i, b_i : i \in L$, that is, omitting index t. It is a fact that the coordinate ring $\mathcal{O}(Z'(r,t))$ is a free $\mathbb{C}[\mathbf{y}]$-module with basis consisting of elements

$$b_t^{m-1} p_{n,L}[e', f'] : \quad m \leq |(\operatorname{supp} e')^c \cap [r] \setminus f'([t-1])|.$$

Digressing briefly, note that the $p_{n,L}[e', f']$ satisfying the above restriction for a given m form a basis of $R(n,L)/I_{n,L}(m,r,t-1)$. Over a general point of $Y_{n,L}(m,r,t-1)$, lying on a component $W_{f'}$, say, there are m choices of $j \in [r] \setminus f'([t-1])$ with $x_j = 0$, and hence of possible y_j values for the extra variable b_t on $Z'(r,t)$ (since a_t must be zero). In a loose sense, this explains why each such $p_{n,L}[e', f']$ should be multiplied by the m different powers b_t^0 through b_t^{m-1} to get a basis of $\mathcal{O}(Z'(r,t))$.

For our given (e, f), we have $n \in (\operatorname{supp} \theta^{-1} e)^c \setminus \theta^{-1} f([t-1])$. Hence if we fix

$$\begin{aligned} m &= |(\operatorname{supp} \theta^{-1} e)^c \cap [n-1] \setminus \theta^{-1} f([t-1])| \\ &= |(\operatorname{supp} \theta^{-1} e)^c \setminus \theta^{-1} f([t-1])| - 1 \\ &= |(\operatorname{supp} e)^c \setminus f^{-1}([t])|, \end{aligned}$$

then the polynomial

$$b_t^m p_{n,L}[\theta^{-1} e, \theta^{-1} f|_L]$$

is one of our basis elements for $\mathcal{O}(Z'(n,t))$ but not for $\mathcal{O}(Z'(n-1,t))$. Note that this is the same m as in Case II of the recurrence for $i(e,f)$. We now define a new polynomial q_1 by combining the above polynomial with a linear combination of our basis elements for $\mathcal{O}(Z'(n-1,t))$ to make it vanish on $Z'(n-1,t)$. In symbols,

$$\begin{aligned} q_1 \in\ & b_t^m p_{n,L}[\theta^{-1} e, \theta^{-1} f|_L] \\ & + \mathbb{C}[\mathbf{y}] \left\{ b_t^{m'-1} p_{n,L}[e', f'] : m' \leq |(\operatorname{supp} e')^c \cap [n-1] \setminus f'([t-1])| \right\}, \\ q_1|_{Z'(n-1,t)} &= 0. \end{aligned}$$

The second condition just means that for each component $V(a_t) \cap W_{f'}$ from the definition of $Z'(n-1,t)$ in (17), the result of substituting $a_i \mapsto x_{f'(i)}$, $b_i \mapsto y_{f'(i)}$, $x_{f'(t)} \mapsto 0$ in q_1 is zero. Thus the above conditions define q_1 in terms of the solution to some linear equations over $\mathbb{C}[\mathbf{y}]$. The existence and uniqueness of the solution is equivalent to the claimed fact that the specified elements of $\mathcal{O}(Z'(n-1,t))$ indeed form a basis, which can be proven using the polygraph theorem for $R(n, l-1)$.

We're not done yet! We now define $q[e,f]$, uniquely modulo $I(1,1,t-1)$, to be a polynomial of the correct bi-degree such that

- $q[e,f]$ coincides with θq_1 on

$$V(x_1) \cap \bigcup_f W_f : \quad f(t) = 1,\ 1 \notin f([t-1]),$$

– $q[e,f] \in I(1,1,t)$, that is, $q[e,f]$ vanishes on

$$V(x_1) \cap \bigcup_f W_f : \quad f(t) \neq 1,\ 1 \notin f([t-1]).$$

Note that the union of the two loci above is $Y(1,1,t-1)$, so the conditions define $q[e,f]$ uniquely modulo $I(1,1,t-1)$, provided a solution exists at all. It is a fact that a solution does exist.

Case III: $t = 0$. We have $e_1 > 0$ and we define

$$q[e,f] = x_1 \theta p[\theta^{-1}(e_1 - 1, e_2, \ldots, e_n), \theta^{-1} f],$$

the right-hand side having already been computed as a basis element of smaller x-degree. In this case, it will turn out that $p[e,f] = q[e,f]$.

Now that the proto-basis elements $q[e,f]$ have been computed, the algorithm concludes with a *correction phase:* define $p[e,f]$ by the conditions

$$\begin{aligned} &p[e,f] \in q[e,f] + \mathbb{C}[\mathbf{y}] \Big\{\ q[e',f'] : t' < t \text{ and} \\ &\qquad\qquad \exists r,k : |(\operatorname{supp} e')^c \cap [r] \setminus f'([k])| > |(\operatorname{supp} e)^c \cap [r] \setminus f([k])|\ \Big\}, \\ &p[e,f] \in I(m,r,k) \text{ for } m > |(\operatorname{supp} e)^c \cap [r] \setminus f([k])|. \end{aligned}$$

In the set of correction terms appearing in the first line above, the condition $t' < t$ refers to the "t" values defined for (e',f') and (e,f) respectively. Thus the correction terms are those $q[e',f']$ defined at a stage with a smaller t value than $q[e,f]$, and for which there is some $I(m,r,k)$ that is supposed to contain $p[e,f]$ but not $p[e',f']$. Note that the correction terms could as well have been $p[e',f']$, presumed defined inductively for smaller t values. Thus the correction serves to force $p[e,f]$ to vanish on the union of its appropriate loci $Y(m,r,k)$, by adjusting it with elements that do not vanish there. It is a fact that this correction is possible, and that the correction terms are linearly independent on the union $\bigcup Y(m,r,k) : m > |(\operatorname{supp} e)^c \cap [r] \setminus f([k])|$, so the correction is unique. As in the construction of q_1 in Case II, the second condition imposes a system of linear equations to be solved for the coefficients of the correction terms.

5.2. ABOUT THE PROOF

Why is it so hard to prove directly that this algorithm gives well-defined results? We have a convenient basis of $\mathbb{C}(\mathbf{y}) \otimes R(n,l)$, namely, the set of rational functions

$$x^e 1_f \in \mathbb{C}(\mathbf{y})[\mathbf{x}, \mathbf{a}, \mathbf{b}]$$

defined in the preceding lecture. The algorithm can be conceived of as a method for computing $p[e,f]$ as a linear combination of the elements $x^e 1_f$.

From this point of view it is actually not so hard to show that it works. The definition of $q[e, f]$ in terms of q_1 in Case II presents no problem in this setting, and one can show that the linear equations over $\mathbb{C}(\mathbf{y})$ defining the correction terms in the construction of q_1 and later of $p[e, f]$ have unique solutions. The problem with this is that it only gives the $p[e, f]$ as elements of $\mathbb{C}(\mathbf{y}) \otimes R(n, l)$. We have no reason to suppose in advance that denominators appearing when we solve various linear equations over $\mathbb{C}(\mathbf{y})$ should cancel out.

Geometrically speaking, working over $\mathbb{C}(\mathbf{y})$ amounts to working with generic values of the $\mathbf{y}$ coordinates. Generically, the $\mathbf{y}$ coordinates are distinct, and the subspaces W_f effectively do not meet. That is precisely why we have the $\mathbb{C}(\mathbf{y})$-basis $\{x^e 1_f\}$. The simplest non-generic points are where exactly two of the $\mathbf{y}$ coordinates coincide, say y_p and y_q, and the rest are distinct from these and each other. At such a point, distinct subspaces W_f and W_g can meet, but only if f and g are close: we can only have $f(i) \neq g(i)$ if $\{f(i), g(i)\} = \{p, q\}$. Locally, in a neighborhood of such a point, $Z(n, l)$ looks like $Z(2, L) \times \mathbb{C}^{2(n-2)}$, for some set of indices $L \subseteq [l]$. This suggests a possible proof strategy. The locus V where there is more than one coincident pair of $\mathbf{y}$ variables has codimension 2 in $\mathbb{C}^n = \operatorname{Spec} \mathbb{C}[\mathbf{y}]$. Any rational function $h(\mathbf{y})$ which is locally regular outside V is regular, *i.e.*, is a polynomial. To show that the linear algebra over $\mathbb{C}(\mathbf{y})$ in the algorithm gives results in $\mathbb{C}[\mathbf{y}]$, it is enough to show that the algorithm works locally outside V, and we can hope to reduce the problem there to the case $n = 2$.

This basic strategy works. Its chief trouble is that it is extremely awkward to formulate the geometric reduction to $n = 2$ directly in terms of the algorithm. This is the main reason why I do not follow an algorithmic approach in [17]. There is a second, more subtle, source of trouble to overcome in the proof, which surfaces in the algorithm in the form of a non-canonical choice. Recall that in Case II, after constructing q_1, we defined $q[e, f]$ by giving its value on $Y(1, 1, t-1)$. The actual polynomial we choose to represent it is arbitrary, modulo $I(1, 1, t-1)$. Part of this arbitrariness is illusory, since different choices might correct to the same $p[e, f]$ in the last phase. But a difference that is not only in $I(1, 1, t-1)$, but also in every $I(m, r, k)$ that contains $p[e, f]$, will persist as a difference in the final $p[e, f]$. In short, the conditions we imposed on the $p[e, f]$ do not define them uniquely.

To make the basis canonical, we can insist in Case II on realizing $q[e, f]$ as a $\mathbb{C}[\mathbf{y}]$-linear combination of basis elements for $R(n, l)/I(1, 1, t-1)$ found by other means (we must resort to other means because we don't yet have the basis of $R(n, l)/I(1, 1, t-1)$ during the proto-basis construction phase). Such elements can be obtained by reverse induction on t using two further facts.

- The elements $q[e,f] = \theta p_{n-1,l}[(\theta^{-1}e)|_{[n-1]}, \theta^{-1}f]$ constructed in Case I form a free $\mathbb{C}[\mathbf{y}]$-module basis of $R(n,l)/I(1,1,l)$.
- Given a free $\mathbb{C}[\mathbf{y}]$-module basis B_t of $R(n,l)/I(1,1,t)$, a free module basis of $R(n,l)/I(1,1,t-1)$ is given by $a_t B_t \cup \{\, b_t^m p_{n,L}[e,f] : m < |(\operatorname{supp} e)^c| \text{ and } p_{n,L}[e,f] \notin I_{n,L}(1,1,t-1) \,\}$, where $L = [l] \setminus \{t\}$.

It is also true that if B_0 is a basis of $R(n,l)/I(1,1,0)$ then $B_0 \cdot \{1, x_1, x_1^2, \ldots\}$ is a basis of $R(n,l)$. These facts provide a much simpler construction of a basis of $R(n,l)$ using inductively the bases $\{p[e,f]\}$ of $R(n-1,l)$ and $R(n,l-1)$. We can't deduce the polygraph theorem from this alone, however, because the simpler basis lacks compatibility with the ideals $I(m,r,k)$, and we need that to make the induction go through.

Exercise: For $n=2$, $l=1$, $e=(0,1)$, $f(1)=1$, use the algorithm to compute the polynomial $p[e,f]$, of bi-degree $(1,0)$. When constructing q_1 in Case II, canonicalize by expressing it in terms of the basis of $R(n,l)/I(1,1,0)$ described above. Your final answer should be $x_1 + x_2 - a_1$, and you should have found $\{a_1, x_2\}$ for the canonicalizing basis of $R(n,l)/I(1,1,0)$ in degree $(1,0)$. Note that $Y(1,2,1)$ is the locus defined by the vanishing of whichever coordinate x_j is not assigned equal to a_1, and that $p[e,f] = x_1 + x_2 - a_1$ vanishes on this locus, as it should.

6. Macdonald polynomials and character formulas

6.1. MACDONALD POLYNOMIALS

In the preceding lectures we've had a fairly thorough look at the ingredients of the proof of the $n!$ theorem, Theorem 1.1, but have so far said nothing about Theorem 1.2, which identifies the doubly-graded character of the space D_μ of derivatives of Δ_μ as a Macdonald polynomial. We haven't even *defined* the Macdonald polynomials, so let's do that now.

We first need the notion of plethystic substitution from the theory of symmetric functions. We work with the algebra of symmetric polynomials

$$\Lambda = \Lambda_{\mathbb{Q}(q,t)}(z_1, z_2, \ldots)$$

in indeterminates z_i, with coefficients in the field of rational functions of two parameters $\mathbb{Q}(q,t)$. In infinitely many indeterminates, the term "symmetric polynomial" is something of a misnomer, since every non-constant symmetric polynomial has infinitely many terms. Strictly speaking, Λ is the algebra of symmetric formal power series with terms of bounded degree. The algebra Λ can be identified with the polynomial ring

$$\Lambda = \mathbb{Q}(q,t)[p_1, p_2, \ldots]$$

in the power-sums

$$p_k(z) = \sum_i z_i^k.$$

For any $A \in \Lambda$, we define

$$p_k[A] = A|_{q \mapsto q^k, t \to t^k, z_i \mapsto z_i^k}$$

to be the result of replacing every indeterminate in A by its k-th power. Since the power-sums freely generate Λ, there is a unique endomorphism of $\mathbb{Q}(q,t)$-algebras

$$\mathrm{ev}_A : \Lambda \to \Lambda; \quad \mathrm{ev}_A(p_k) = p_k[A].$$

We define for every $f \in \Lambda$ the *plethystic substitution*

$$f[A] = \mathrm{ev}_A(f).$$

We also set $Z = p_1 = z_1 + z_2 + \cdots$, and we see immediately that $p_k[Z] = p_k(z)$, hence

$$f[Z] = f(z)$$

for all $f \in \Lambda$. More generally, if $A = a_1 + a_2 + \cdots$ is any sum of monomials, each with coefficient 1, then $p_k[A] = p_k(a_1, a_2, \ldots)$ and hence $f[A] = f(a_1, a_2, \ldots)$ for all f. The only plethystic substitutions we need here are

$$f \mapsto f[Z/(1-q)], \quad f \mapsto f[Z/(1-t)],$$

and their inverses

$$f \mapsto f[Z(1-q)], \quad f \mapsto f[Z(1-t)].$$

If we identify $\mathbb{Q}(q,t)$ with a subfield of the formal Laurent series in q, t, we can conveniently interpret $f[Z/(1-q)]$ as

$$f[Z(1+q+q^2+\cdots)] = f(z_1, z_2, \ldots, qz_1, qz_2, \ldots, q^2 z_1, q^2 z_2, \ldots).$$

THEOREM/DEFINITION 6.1. *There is a homogeneous $\mathbb{Q}(q,t)$-basis $\{\tilde{H}_\mu(z)\}$ of Λ whose elements are uniquely characterized by the conditions*

(1) $\tilde{H}_\mu[Z(1-q)] \in \mathbb{Q}(q,t)\{s_\lambda : \lambda \geq \mu\}$,
(2) $\tilde{H}_\mu[Z(1-t)] \in \mathbb{Q}(q,t)\{s_\lambda : \lambda \geq \mu'\}$, *and*
(3) $\tilde{H}_\mu[1] = 1$.

Here μ is an integer partition, and λ ranges over partitions of the same integer $|\lambda| = |\mu|$, with $\leq$ denoting the dominance partial ordering on partitions.

The above theorem is a recasting of Macdonald's existence theorem. The $\tilde{H}_\mu$ are related to the integral form Macdonald polynomials J_μ, as defined in Macdonald's book [24] , by

$$\tilde{H}_\mu(z) = t^{n(\mu)} J_\mu \left[\frac{Z}{1-t^{-1}}; q, t^{-1} \right].$$

Equivalently, if we define coefficients $\tilde{K}_{\lambda\mu}(q,t) \in \mathbb{Q}(q,t)$ through the Schur function expansion

$$\tilde{H}_\mu(z) = \sum_\lambda \tilde{K}_{\lambda\mu} s_\lambda(z),$$

then these are related to Macdonald's q,t-Kostka coefficients $K_{\lambda\mu}(q,t)$ by

$$\tilde{K}_{\lambda\mu}(q,t) = t^{n(\mu)} K_{\lambda\mu}(q, t^{-1}).$$

Note that since $n(\mu)$ is the x-degree of Δ_μ, Theorem 1.2 shows that $\tilde{K}_{\lambda\mu}(q,t)$ not only is a polynomial with coefficients in $\mathbb{N}$, but that its degree in t is $n(\mu)$. It follows that $K_{\lambda\mu}(q,t)$ is also a polynomial with coefficients in $\mathbb{N}$, which was Macdonald's original conjecture.

Theorem 6.1 is deduced from Macdonald's existence theorem in [24] as follows. Condition (2) is the triangularity condition for Macdonald's J_μ. Condition (1) follows from the symmetry $\tilde{H}_{\mu'}(z;q,t) = \tilde{H}_\mu(z;t,q)$, or $\tilde{K}_{\lambda,\mu'}(q,t) = \tilde{K}_{\lambda,\mu}(t,q)$, a consequence of the symmetries (8.14–15) in Macdonald's book. Condition (3) says $\tilde{K}_{(n),\mu} = 1$, which follows from formula (6.11′) in Macdonald's book with $n \to \infty$ in his formula. This shows that the basis $\{\tilde{H}_\mu(z)\}$ exists. It is unique because conditions (1) and (2) imply that any alternative basis must be both upper- and lower-triangular with respect to $\{\tilde{H}_\mu(z)\}$ in the dominance ordering. Hence any alternative would consist of scalar multiples of the $\tilde{H}_\mu(z)$, and condition (3) fixes the scalar factor. We remark that the orthogonality property in Macdonald's original definition can be recovered from conditions (1) and (2) and the orthogonality of Schur functions with respect to the usual Hall inner product.

6.2. EXAMPLE

We compute $\tilde{H}_\mu(z)$ for all $|\mu| = 3$. First, we find the plethystically transformed Schur functions

$$\begin{aligned} s_3[Z/(1-q)] &= \tfrac{1}{(1-q)(1-q^2)(1-q^3)} \left[s_3(z) + (q+q^2) s_{21}(z) + q^3 s_{111} \right], \\ s_{21}[Z/(1-q)] &= \tfrac{1}{(1-q)(1-q^2)(1-q^3)} \\ &\quad \cdot \left[(q+q^2) s_3(z) + (1+q+q^2+q^3) s_{21}(z) \right. \\ &\quad\quad \left. + (q+q^2) s_{111}(z) \right]. \end{aligned}$$

By condition (1) in the definition, $\tilde{H}_3(z)$ is a scalar multiple of $s_3[Z/(1-q)]$. Condition (3) determines the scalar:

$$\tilde{H}_3(z) = s_3(z) + (q+q^2)s_{21}(z) + q^3 s_{111}.$$

By symmetry we also have

$$\tilde{H}_{111}(z) = s_3(z) + (t+t^2)s_{21}(z) + t^3 s_{111}.$$

By condition (1), $\tilde{H}_{21}(z)$ is a linear combination of $s_{21}[Z/(1-q)]$ and $s_3[Z/(1-q)]$. Note that $s_{21}[Z/(1-q)] - (q+q^2)s_3[Z/(1-q)]$ is a scalar multiple of $s_3(z) + qs_{21}(z)$, while $s_{21}[Z/(1-q)] - (q^{-1}+q^{-2})s_3[Z/(1-q)]$ is a scalar multiple of $s_{21}(z) + qs_{111}(z)$. Therefore

$$\tilde{H}_{21}(z) = s_3(z) + (q+t)s_{21}(z) + qts_{111}(z)$$

satisfies condition (1), and hence also condition (2), by symmetry. Clearly it also satisfies condition (3).

Let's compare this with Theorem 1.2 for $\mu = (21)$. The polynomial

$$\Delta_{21}(\mathbf{x}, \mathbf{y}) = x_1y_2 - x_1y_3 + x_2y_3 - x_2y_1 + x_3y_1 - x_3y_2$$

has degree $(1,1)$ and is S_n-alternating, so the space of derivatives D_{21} has the sign character χ^{111} in degree $(1,1)$. This accounts for the term $qts_{111}(z)$ in $\tilde{H}_{21}(z)$. Differentiating by the $\mathbf{y}$ variables gives the linear span of $x_1 - x_2$, $x_2 - x_3$ in degree $(1,0)$, with character χ^{21}. Similarly $y_1 - y_2$, $y_2 - y_3$ span degree $(0,1)$ with the same character. These two degrees account for the term $(q+t)s_{21}(z)$. Finally, the constants in degree $(0,0)$ have the trivial character χ^3, accounting for the term $s_3(z)$.

6.3. FROBENIUS SERIES

To prove Theorem 1.2 using our geometric results, the first step is to recast the theorem in more direct symmetric function terms by introducing the *Frobenius characteristic*. This is the map

$$\Phi\colon (S_n \text{ characters}) \to \Lambda^{(n)},$$

where $\Lambda^{(n)}$ denotes the space of symmetric polynomials homogeneous of degree n, given by

$$\Phi\chi = \frac{1}{n!} \sum_{w \in S_n} \chi(w) p_{\tau(w)}(z).$$

Here $\tau(w)$ is the partition of n whose parts are the cycle-lengths in the expression for w as a product of disjoint cycles. By a theorem of Frobenius, the irreducible characters χ^λ of the symmetric groups are given in terms of symmetric functions by the identity

$$\Phi\chi^\lambda = s_\lambda(z).$$

Using this, we can reformulate Theorem 1.2 as an identity

$$\mathcal{F}_{D_\mu}(z;q,t) \underset{\text{def}}{=} \sum_{r,s} t^r q^s \Phi \mathrm{ch}(D_\mu)_{r,s} = \tilde{H}_\mu(z). \tag{18}$$

We call the expression $\mathcal{F}_{D_\mu}(z;q,t)$ the *Frobenius series* of the doubly-graded S_n-module D_μ.

Now D_μ is isomorphic as a doubly-graded S_n module to $R_\mu = \mathbb{C}[\mathbf{x},\mathbf{y}]/J_\mu$, and by Theorem 3.4, the latter is the coordinate ring of the scheme-theoretic fiber of X_n over the distinguished point $I_\mu \in H_n$. Equivalently, R_μ is the fiber of the vector bundle P on H_n at I_μ. We will prove that $\mathcal{F}_{R_\mu}(z;q,t)$ satisfies the conditions (1)–(3) characterizing $\tilde{H}_\mu(z)$ in Theorem 6.1. To do this, we extend the notion of Frobenius series to a broader geometric context. A fuller discussion is given in [16].

The algebraic torus

$$\mathbb{T}^2 = (\mathbb{C}^*)^2$$

acts on $\mathbb{C}^2$ as the group of 2×2 invertible diagonal matrices. This action extends equivariantly to all geometric objects under consideration: in particular, to the schemes H_n and X_n, the universal family F, the polygraph $Z(n,l)$, the vector bundles B and P. The monomial ideals $I_\mu \in H_n$ are precisely the $\mathbb{T}^2$-fixed points of H_n. Wherever we have encountered doubly graded algebraic objects, the double grading can be identified with the $\mathbb{T}^2$ action, in the sense that an element p is doubly homogeneous of degree (r,s) if and only if p is a simultaneous eigenvector for the $\mathbb{T}^2$ action, with $(t,q) \in \mathbb{T}^2$ acting by

$$(t,q)\cdot p = t^r q^s p.$$

For example, an ideal $I \subseteq \mathbb{C}[x,y]$ is $\mathbb{T}^2$-invariant if and only if it is doubly homogeneous—and hence a monomial ideal. For another example, the double grading of R_μ reflects the $\mathbb{T}^2$ action on the fiber of the $\mathbb{T}^2$-equivariant vector bundle P over the $\mathbb{T}^2$-fixed point I_μ. From this point of view, the Frobenius series of a doubly-graded S_n-module should be regarded as a formal expression for its character as a $\mathbb{T}^2 \times S_n$-module.

Consider the local ring of the Hilbert scheme at a $\mathbb{T}^2$-fixed point,

$$S = \mathcal{O}_{H_n, I_\mu},$$

with maximal ideal M. The local ring (S,M) is regular, and has a $\mathbb{T}^2$ action. Since I_μ is an isolated fixed point, the induced action of $\mathbb{T}^2$ on the cotangent space M/M^2 has no non-trivial invariant vectors. Therefore the rational function $\det_{M/M^2}(1-\tau)$ of $\tau = (t,q) \in \mathbb{T}^2$ is non-zero. We can make this considerably more precise. For every n-element subset $L \subseteq \mathbb{N} \times \mathbb{N}$, we have seen that

$$\Delta_L/\Delta_\mu$$

represents a regular function on H_n in a neighborhood of I_μ. Provided L is not equal to the diagram of μ, this function vanishes at I_μ, that is, it belongs to the maximal ideal M of S. In [15] we gave an alternate proof of Fogarty's theorem by exhibiting a specific system of regular local parameters (spanning set of M/M^2) consisting of $2n$ of the the above elements, selected as follows. Choose a cell x in the diagram of μ, and locate the cells just inside and outside the ends of the hook based at x, as indicated in the example below.

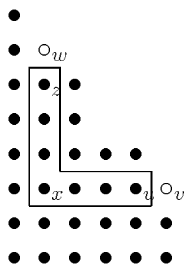

Here the solid dots represent cells in the diagram of μ, while the open dots are just outside. This choice contributes the two elements

$$\Delta_{d(\mu)\cup\{v\}\setminus\{z\}}/\Delta_\mu, \quad \Delta_{d(\mu)\cup\{w\}\setminus\{u\}}/\Delta_\mu$$

to the regular system of parameters. Their respective degrees are

$$(-l, 1+a), \quad (1+l, -a),$$

where a and l are the lengths of the *arm* and *leg* in the hook of x (not counting x itself). These parameters form a basis of $\mathbb{T}^2$-eigenvectors in M/M^2. Hence we have explicitly

$$\det_{M/M^2}(1-\tau_{t,q}) = \prod_{x\in\mu}(1-t^{-l(x)}q^{1+a(x)})(1-t^{1+l(x)}q^{-a*x)}). \qquad (19)$$

This formula was first obtained by Ellingsrud and Strömme [5], in a somewhat different equivalent form.

For any $\mathbb{T}^2$-equivariant finitely generated S-module T, the modules $\mathrm{Tor}_i^S(S/M,T)$ are finite-dimensional doubly graded vector spaces. If T also has an action of S_n by S-module automorphisms, then these Tor-modules are S_n modules as well, and so have Frobenius series, with coefficients that are Laurent polynomials in q and t. We now define

$$\mathcal{F}_T(z;q,t) = \frac{\sum_i (-1)^i \mathcal{F}_{\mathrm{Tor}_i^S(S/M,T)}(z;q,t)}{\det_{M/M^2}(1-\tau_{t,q})}.$$

The sum is finite, since S is regular and therefore T has a finite free resolution. The long exact sequence for Tor shows that $\mathcal{F}_T$ is additive on exact sequences:

$$\mathcal{F}_T = \mathcal{F}_{T'} + \mathcal{F}_{T''} \quad \text{for} \quad 0 \to T' \to T \to T'' \to 0.$$

Using the Koszul resolution of S/M, one sees that for $T = V^\lambda$, an irreducible character of S_n, with $\mathbb{T}^2$ acting by some character $(t,q)\cdot v = t^r q^s v$, and trivial S-module structure $V^\lambda = (S/M)\otimes_{\mathbb{C}} V^\lambda$, our new definition of $\mathcal{F}_T$ yields $t^r q^s s_\lambda(z)$, in agreement with the Frobenius series as defined before. By additivity, it follows that our generalized Frobenius "series" coincides with the original whenever T has finite length, *i.e.*, is finite-dimensional over $\mathbb{C}$.

6.4. THE DOUBLY-GRADED CHARACTER OF D_μ

The key result we need for Theorem 1.2 is as follows.

THEOREM 6.2. *Let T be a finite S-algebra with an equivariant $\mathbb{T}^2$-action and an S_n-action commuting with the $\mathbb{T}^2$-action. Suppose that $y_1,\dots,y_n \in T$ are doubly homogeneous of degree (r,s) and form a regular sequence in T, and that S_n acts by permuting them. Then we have*

$$\mathcal{F}_{T/(\mathbf{y})T}(z;q,t) = \mathcal{F}_T[Z(1-t^r q^s);q,t].$$

This theorem is proved using the Koszul resolution of $T/(\mathbf{y})T$ as a T-module, together with the following property of the Frobenius characteristic: if E_k denotes the k-th exterior power of the natural representation of S_n, and V is any representation, then

$$\sum_k (-u)^k \Phi\mathrm{ch}(E_k \otimes V) = f[Z(1-u)], \quad \text{where } f(z) = \Phi\mathrm{ch}(V).$$

Now let's see how to prove that $\mathcal{F}_{R_\mu}(z;q,t) = \tilde{H}_\mu(z)$. We take $S = \mathcal{O}_{H_n,I_\mu}$ as before and $T = \mathcal{O}_{X_n,Q_\mu}$, where $Q_\mu \in X_n$ is the unique point lying above I_μ. By Theorem 3.4 we know that T is a free S-module, with basis any $\mathbb{C}$-basis of $R_\mu = T/MT$. In particular, we can choose such a basis to be doubly homogeneous and span an S_n-invariant space of polynomials.

Thus we have an equivariant isomorphism of S-modules $T \cong R_\mu \otimes_{\mathbb{C}} S$. It follows that

$$\mathcal{F}_T(z;q,t) = \frac{\mathcal{F}_{R_\mu}(z;q,t)}{\det_{M/M^2}(1-\tau_{t,q})}.$$

We also know that $\mathbf{y} = y_1,\ldots,y_n$ is a regular sequence in T, as a consequence of Lemma 4.1. Theorem 6.2 yields

$$\mathcal{F}_{T/(\mathbf{y})T}(z;q,t) = \mathcal{F}_T[Z(1-q);q,t].$$

Combining the above two equations, we obtain condition (1) in Theorem 6.1 for $\mathcal{F}_{R_\mu}(z;q,t)$, provided that

$$\mathcal{F}_{T/(\mathbf{y})T}(z;q,t) \in \mathbb{Q}(q,t)\{s_\lambda(z) : \lambda \geq \mu\}.$$

For this it suffices to show that there exists a $\mathbb{T}^2 \times S_n$-invariant subspace $V_\mu \subseteq T/(\mathbf{y})T$, which generates $T/(\mathbf{y})T$ as an S-module, such that only the characters $\chi^\lambda : \lambda \geq \mu$ occur with non-zero multiplicity in $\mathrm{ch}(V_\mu)$. A suitable space V_μ is $R_\mu/(\mathbf{y})R_\mu$. Since R_μ generates T, $R_\mu/(\mathbf{y})R_\mu$ generates $T/(\mathbf{y})T$ as an S-module. Since we killed off the $\mathbf{y}$ variables, $R_\mu/(\mathbf{y})R_\mu$ is a quotient of the polynomial ring $\mathbb{C}[\mathbf{x}]$. It is isomorphic as a (singly) graded S_n-module to the y-degree zero component $(D_\mu)_{-,0}$ of D_μ. The latter space is the Macaulay inverse system generated by the *Garnir polynomials*, the derivatives of Δ_μ by maximal degree monomials in $\mathbf{y}$. The Garnir polynomials are the product of Vandermonde determinants

$$\Delta(x_1,\ldots,x_{\mu'_1})\Delta(x_{\mu'_1+1},\ldots,x_{\mu'_1+\mu'_2})\cdots\Delta(x_{n-\mu'_l+1,\ldots,x_n})$$

and its conjugates by the S_n-action permuting the variables. They span the unique copy of the irreducible representation V^μ of minimal degree, namely $n(\mu)$, in the polynomial ring $\mathbb{C}[\mathbf{x}]$. The Macaulay inverse system they generate was studied by Garsia, Procesi and N. Bergeron [2, 11], who showed by elementary means that its graded character is given by the classical t-Kostka coefficients

$$\sum_r t^r\langle\chi^\lambda, \mathrm{ch}(D_\mu)_{r,0}\rangle = t^{n(\mu)}K_{\lambda\mu}(t^{-1}) = \tilde{K}_{\lambda\mu}(0,t).$$

This result is what I alluded to in Lecture 1 as the starting point for my work with Garsia on the $n!$ conjecture. It yields the simplest known proof of the positivity theorem for the classical coefficients $K_{\lambda\mu}(t)$. Since it is known that $K_{\lambda\mu}(t) = 0$ unless $\lambda \geq \mu$, we have established condition (1) for $\mathcal{F}_{R_\mu}(z;q,t)$. Condition (2) follows by symmetry. Condition (3) just says that the unique invariants in R_μ are the constants, which is clear. Theorem 1.2 follows.

6.5. MORE CHARACTER FORMULAS

Let A be any $\mathbb{T}^2$-equivariant coherent sheaf of $\mathcal{O}_{H_n}$-modules. Since the Chow morphism $H_n \to S^n\mathbb{C}^2$ is projective, the cohomology modules

$$H^i(H_n, A)$$

are finitely generated doubly-graded modules over $\mathcal{O}(S^n\mathbb{C}^2) = \mathbb{C}[\mathbf{x}, \mathbf{y}]^{S_n}$. Their Hilbert series

$$\mathcal{H}^i_A(q,t) \underset{\text{def}}{=} \mathcal{H}_{H^i(H_n,A)}(q,t)$$

are formal Laurent series which are rational functions of q and t. If A has an S_n-action we can similarly define Frobenius series

$$\mathcal{F}^i_A(z;q,t) \underset{\text{def}}{=} \mathcal{F}_{H^i(H_n,A)}(z;q,t) \in \Lambda.$$

The cohomology Euler characteristic of A is given by the *Atiyah–Bott formula*

$$\sum_i (-1)^i \mathcal{H}^i_A(q,t) = \sum_\mu \frac{\sum_i (-1)^i \mathcal{H}_{\mathrm{Tor}_i(k_\mu, A)}(q,t)}{\det_{T^*H_n(I_\mu)}(1 - \tau_{t,q})}.$$

Here $k_\mu = S/M$ is the residue field of $S = \mathcal{O}_{H_\mu, I_\mu}$, regarded as a "skyscraper sheaf" supported at the point I_μ, and T^*H_n denotes the cotangent bundle, so $T^*H_n(I_\mu) = M/M^2$, where M is the maximal ideal of S. If A has an S_n-action, then we can apply the Atiyah–Bott formula to each isotypic component of A separately, obtaining

$$\sum_i (-1)^i \mathcal{F}^i_A(q,t) = \sum_\mu \frac{\sum_i (-1)^i \mathcal{F}_{\mathrm{Tor}_i(k_\mu, A)}(q,t)}{\det_{T^*H_n(I_\mu)}(1 - \tau_{t,q})}.$$

Note that the term on the right for any particular μ is nothing other than our definition of $\mathcal{F}_{A_{I_\mu}}(z;q,t)$ for the stalk A_{I_μ} as a module over the local ring S. Our context differs from the original setting for the Atiyah–Bott formula in that H_n is not projective. However, there are more general versions of the formula which apply in our situation [27].

When A is suitably explicit, the right-hand side of the Atiyah–Bott formula can be written out in full. We already know the denominator, which is given by (19). If A is locally free, that is, if A is a vector bundle, then the only non-zero term in the numerator is the Hilbert or Frobenius series of the fiber $A(I_\mu)$. For example, we have seen in Lecture 4 that there is a natural homomorphism

$$R(n,l) \to H^0(X_n, B^{\otimes l}) = H^0(H_n, P \otimes B^{\otimes l}).$$

The second equality here results from the definition $P = \rho_* \mathcal{O}_{X_n}$. We remarked earlier and will soon show that this homomorphism is an isomorphism and that the higher cohomology modules $H^i(H_n, P \otimes B^{\otimes l})$ vanish for $i > 0$. Taking $A = P \otimes B^{\otimes l}$, the left-hand side of the Atiyah–Bott formula thus reduces to the Frobenius series $\mathcal{F}_{R(n,l)}(z; q, t)$, where S_n acts on $R(n, l)$ through the variables $\mathbf{x}, \mathbf{y}$. For the right-hand side, we have

$$\mathcal{F}_{P(I_\mu)}(z; q, t) = \mathcal{F}_{R_\mu}(z; q, t) = \tilde{H}_\mu(z),$$

and we have

$$\mathcal{H}_{B(I_\mu)}(q, t) = \mathcal{H}_{\mathbb{C}[x,y]/I_\mu}(q, t) = B_\mu(q, t) \underset{\text{def}}{=} \sum_{(r,s)\in\mu} t^r q^s.$$

Therefore we obtain the explicit formula

$$\mathcal{F}_{R(n,l)}(z; q, t) = \sum_\mu \frac{B_\mu(q, t)^l \tilde{H}_\mu(z)}{\prod_{x\in\mu}(1 - t^{-l(x)} q^{1+a(x)})(1 - t^{1+l(x)} q^{-a(x)})}.$$

To simplify this, we introduce the linear operator

$$\Delta : \Lambda \to \Lambda; \quad \Delta \tilde{H}_\mu(z) \underset{\text{def}}{=} B_\mu(q, t) \tilde{H}_\mu(z).$$

We also observe that for $l = 0$ we have $R(n, 0) = \mathbb{C}[\mathbf{x}, \mathbf{y}]$, and the Frobenius series of this is given in plethystic notation by

$$\mathcal{F}_{\mathbb{C}[\mathbf{x},\mathbf{y}]}(z; q, t) = h_n\left[\frac{Z}{(1-q)(1-t)}\right],$$

where $h_n(z) = s_{(n)}(z)$ is the complete homogeneous symmetric function. Hence we obtain

$$\mathcal{F}_{R(n,l)}(z; q, t) = \Delta^l h_n\left[\frac{Z}{(1-q)(1-t)}\right].$$

We can specialize any Frobenius series $\mathcal{F}$ to the corresponding Hilbert series by specializing the Schur function $s_\lambda(z)$ to $f_\lambda = \chi^\lambda(1) = \dim V^\lambda$, the number of standard Young tableaux of shape λ. By the formula for the Frobenius characteristic, the result is equal to $n!$ times the coefficient of p_1^n in the power-sum expansion of $\mathcal{F}$, or to $\partial p_1^n \mathcal{F}$, if we regard $\mathcal{F}$ as a polynomial in the power-sums. This gives a second formula for the Hilbert series of $R(n, l)$, which we had earlier determined combinatorially:

$$\begin{aligned} \mathcal{H}_{R(n,l)}(q, t) &= A^{n,l}(q, t) = \frac{1}{(1-q)^n} \sum_{e,f} t^{|e|} q^{i(e,f)} \\ &= \partial p_1^n \Delta^l h_n\left[\frac{Z}{(1-q)(1-t)}\right]. \end{aligned}$$

Note that the last equation is between a purely combinatorial quantity and a quantity computed using only symmetric function theory. I leave it as a challenge to you to try to prove it by elementary means!

7. The diagonal harmonics character formula

7.1. MORE CHARACTER FORMULAS

In the last lecture, we saw how to use the Atiyah–Bott theorem to write down a formula for the Frobenius series of $R(n,l)$. For this we had to assume that the natural map

$$R(n,l) \to H^0(H_n, P \otimes B^{\otimes l}) \tag{20}$$

is an isomorphism and that the higher sheaf cohomology groups $H^i(H_n, P \otimes B^{\otimes l})$, $i > 0$, vanish. We'll discuss how to justify this assumption at the end of the lecture. First we'll investigate how to get various other, related formulas.

The Frobenius series $\mathcal{F}_{R(n,l)}$ contains information about $R(n,l)$ as a doubly graded S_n-module. But $R(n,l)$ is actually a doubly graded $(S_n \times S_l)$-module, and we might ask to describe it as such. Under the identification in (20), the S_n action comes from P and the S_l action comes from permutations of the factors in $B^{\otimes l}$. The decomposition into S_l-isotypic components is given by Schur-Weyl duality as

$$B^{\otimes l} = \bigoplus_{\nu} S_\nu(B) \otimes V^\nu,$$

where ν ranges over partitions of l, V^ν denotes an irreducible S_l-module, and S_ν is the corresponding *Schur functor*. The Schur functor is characterized by the fact that the representation of $\mathrm{GL}(W)$ on $S_\nu(W)$ is the irreducible representation with character given by the Schur function $s_\nu(z)$. More precisely, if $\tau : W \to W$ is a diagonalizable endomorphism with eigenvalues $z_1, \ldots, z_k$, then the induced action of τ on $S_\nu(W)$ has trace equal to $s_\nu(z_1, \ldots, z_k)$. It follows that

$$\mathrm{Hom}_{S_l}(V^\nu, R(n,l)) \cong H^0(H_n, P \otimes S_\nu(B)),$$

and that the higher cohomologies vanish. The doubly graded Hilbert series of $S_\nu(B(I_\mu))$, or equivalently, the trace of $\tau_{t,q} \in \mathbb{T}^2$ on $S_\nu(B(I_\mu))$, is equal to the Schur function s_ν evaluated on the eigenvalues of $\tau_{t,q}$ on $B(I_\mu)$. The sum of these eigenvalues is the generating function for the cells in the diagram of μ, which we previously denoted by

$$B_\mu(q,t) = \sum_{(r,s) \in \mu} t^r q^s.$$

It is convenient to express the Schur function of the eigenvalues in plethystic notation as

$$\mathcal{H}_{S_\nu(B(I_\mu))}(q,t) = S_\nu[B_\mu(q,t)].$$

Setting

$$R(n,l)^\nu = \mathrm{Hom}_{S_l}(V^\nu, R(n,l)),$$

its Frobenius series is now given by the Atiyah–Bott formula as

$$\mathcal{F}_{R(n,l)^\nu}(z;q,t) = \sum_\mu \frac{S_\nu[B_\mu(q,t)]\tilde{H}_\mu(z)}{\prod_{x\in\mu}(1-t^{-l(x)}q^{1+a(x)})(1-t^{1+l(x)}q^{-a(x)})}.$$

Again we can simplify this by introducing the operator

$$\nabla_\nu : \Lambda \to \Lambda, \quad \nabla_\nu \tilde{H}_\mu(z) = S_\nu[B_\mu(q,t)]\tilde{H}_\mu(z).$$

In this notation, the operator Δ introduced earlier is $\nabla_{(1)}$. Note that products of the operators ∇_ν are positive integral linear combinations of these operators, since products of Schur functions have the same property. In particular,

$$\Delta^l = \nabla_{(1)}^l = \sum_{|\nu|=l} f_\nu \nabla_\nu,$$

where $f_\nu = \chi^\nu(1)$ is the number of standard Young tableaux of shape ν. Taking account of the S_l action, our earlier identity $\mathcal{F}_{R(n,l)} = \Delta^l h_n[Z/(1-q)(1-t)]$ is refined to the system of identities

$$\mathcal{F}_{R(n,l)^\nu}(z;q,t) = \nabla_\nu h_n\left[\frac{Z}{(1-q)(1-t)}\right].$$

For every ν, the expression above is "q,t-Schur positive," that is, it is a linear combination of Schur functions by power series in q and t with non-negative integer coefficients. Our present knowledge of $R(n,l)$ is not adequate to provide a combinatorial interpretation of the above expressions, or even of $\mathcal{F}_{R(n,l)}$. What we lack is a basis of $R(n,l)$ compatible with the S_n and S_l actions.

7.2. THE ZERO-FIBER IN THE HILBERT SCHEME

Let

$$Z_n = \sigma^{-1}(\{\underline{0}\}) \subseteq H_n$$

denote the fiber of the Chow morphism $\sigma\colon H_n \to S^n\mathbb{C}^2$ over the origin $\underline{0} \in S^n\mathbb{C}^2$. Here we mean the set-theoretic zero-fiber, as a reduced subscheme of H_n. In fact there is no ambiguity, since it was shown in [15] that the scheme-theoretic zero-fiber is reduced. Our next step is to work out character formulas for the vector bundles $P \otimes B^{\otimes l}$ restricted to the zero-fiber. As before, we will explicitly identify the spaces of global sections

$$H^0(Z_n, P \otimes B^{\otimes l}),$$

show that the higher cohomologies vanish, and use the Atiyah–Bott formula to express the Frobenius series. I'll do this in outline, referring you to [18] for the details.

The key observation is that $\mathcal{O}_{Z_n}$ has an explicit $\mathbb{T}^2$-equivariant resolution by locally free sheaves of $\mathcal{O}_{H_n}$-modules. To write it down we need a bit of notation. Let $\mathbb{C}_t$ and $\mathbb{C}_q$ denote the 1-dimensional representations of $\mathbb{T}^2$ on which $\tau_{t,q} \in \mathbb{T}^2$ acts by t and q, respectively. We write

$$\mathcal{O}_t = \mathcal{O}_{H_n} \otimes \mathbb{C}_t, \quad \mathcal{O}_q = \mathcal{O}_{H_n} \otimes \mathbb{C}_q$$

for $\mathcal{O}_{H_n}$ with its natural $\mathbb{T}^2$ action twisted by these 1-dimensional characters. You can think of them as copies of $\mathcal{O}$ with degree shifts of $(1, 0)$ and $(0, 1)$, respectively. There is a *trace homomorphism* of $\mathcal{O}_{H_n}$-modules

$$\mathrm{tr}: B \to \mathcal{O}$$

defined as follows. Let $\alpha \in B(U)$ be a section of B on some open set U. Since B is a sheaf of algebras and also a vector bundle, there is a regular function $\mathrm{tr}\,(\alpha) \in \mathcal{O}(U)$ whose value at I is the trace of multiplication by α on the fiber $B(I)$. Explicitly, B is a quotient of the trivial bundle $\mathbb{C}[x, y] \otimes \mathcal{O}$, and we have

$$\mathrm{tr}\,(x^r y^s) = p_{r,s}(\mathbf{x}, \mathbf{y}). \tag{21}$$

Note that every S_n-invariant polynomial in $\mathbf{x}, \mathbf{y}$ defines a global regular function on $S^n\mathbb{C}^2$, and hence on H_n, via the Chow morphism. The validity of (21) is easily seen by considering the dense open set of points $I \in H_n$ where $\sigma(I)$ is a set of n distinct points $\{(x_1, y_1), \ldots, (x_n, y_n)\}$. From equation (21), we see that

$$\frac{1}{n}\mathrm{tr}: B \to \mathcal{O}$$

is left inverse to the canonical inclusion

$$\mathcal{O} \hookrightarrow B,$$

so we have a direct-sum decomposition of $\mathcal{O}_{H_n}$-module sheaves, or of vector bundles,

$$B = \mathcal{O} \oplus B', \quad \text{where} \quad B' = \ker(\mathrm{tr}).$$

THEOREM 7.1. *There is a canonical $\mathbb{T}^2$-equivariant resolution*

$$\cdots \to B \otimes \wedge^k(B' \oplus \mathcal{O}_t \oplus \mathcal{O}_q) \to \cdots \to B \otimes (B' \oplus \mathcal{O}_t \oplus \mathcal{O}_q) \to B \to \mathcal{O}_{Z_n} \to 0. \tag{22}$$

Proof. (sketch) The projection $\pi\colon F \to H_n$ of the universal family on H_n is set-theoretically one-to-one over Z_n, so induces a bijective morphism of the reduced preimage $\tilde{Z}_n = \pi^{-1}(Z_n)$ onto Z_n. This morphism corresponds to a homomorphism $\mathcal{O}_{Z_n} \to \mathcal{O}_{\tilde{Z}_n}$. Since $B = \pi_* \mathcal{O}_F$ is generated over $\mathcal{O}_{H_n}$ by x and y, which vanish identically on $\tilde{Z}_n$, the homomorphisms $\mathcal{O}_{H_n} \to \mathcal{O}_{Z_n} \to \mathcal{O}_{\tilde{Z}_n}$ and therefore $\mathcal{O}_{Z_n} \to \mathcal{O}_{\tilde{Z}_n}$ are surjective. This shows that $\tilde{Z}_n \to Z_n$ is a closed embedding and hence an isomorphism. The inclusion $\tilde{Z}_n \subseteq F$ corresponds to a surjective homomorphism $B \to \mathcal{O}_{\tilde{Z}_n} \cong \mathcal{O}_{Z_n}$. This gives the final map in the resolution (22).

The kernel of this map is an ideal in B containing the sections x and y. Multiplication by x and y respectively gives maps $B \otimes \mathcal{O}_t \to B$, $B \otimes \mathcal{O}_q \to B$ with image in the kernel of $B \to \mathcal{O}_{Z_n}$. The subbundle $B' = \ker(\mathrm{tr}) \subseteq B$ is generated by the sections

$$x^r y^s - \frac{1}{n} p_{r,s}(\mathbf{x}, \mathbf{y}),$$

and these also belong to the kernel of $B \to \mathcal{O}_{Z_n}$. For $r + s > 0$ this is clear, since x, y, and $p_{r,s}(\mathbf{x}, \mathbf{y})$ all vanish on Z_n. For $r = s = 0$ the displayed expression is identically zero. Multiplication therefore gives a map $B \otimes B' \to B$ with image in the kernel of $B \to \mathcal{O}_{Z_n}$. Combining these gives the second-to-last map in (22).

Now F is flat and finite over the non-singular variety H_n, and hence Cohen-Macaulay. By an old theorem of Briançon, the zero fiber Z_n is irreducible of dimension $n - 1$, or codimension $n + 1$ in F. Note that $n + 1$ is also the rank of the vector bundle $B' \oplus \mathcal{O}_t \oplus \mathcal{O}_q$, so that the image of the map

$$B \otimes (B' \oplus \mathcal{O}_t \oplus \mathcal{O}_q) \to B$$

is a sheaf of ideals locally generated by $n + 1$ elements. One shows as in [15] that this sheaf of ideals defines Z_n set-theoretically, and generically also scheme-theoretically. Then it is the ideal sheaf of a generically reduced local complete intersection in F, which is therefore reduced and equal to Z_n. The whole sequence in (22) is now exact, since it is the Koszul resolution of Z_n as a local complete intersection in F. As a corollary we can also deduce, as remarked above, that the scheme theoretic zero-fiber is reduced. □

The higher cohomology vanishing $H^i(H_n, P \otimes B^{\otimes l}) = 0$ for $i > 0$ implies a similar vanishing for the tensor product of every term in (22) with P or with $P \otimes B^{\otimes l}$. This immediately implies that

$$H^i(Z_n, P \otimes B^{\otimes l}) = H^i(H_n, \mathcal{O}_{Z_n} \otimes P \otimes B^{\otimes l}) = 0 \quad \text{for } i > 0$$

and also that

$$R(n, l+1) \cong H^0(H_n, B \otimes P \otimes B^{\otimes l}) \to H^0(Z_n, P \otimes B^{\otimes l})$$

is surjective. It also determines the kernel of this surjection as the image of the map induced by

$$B \otimes (B' \oplus \mathcal{O}_t \oplus \mathcal{O}_q) \to B,$$

tensored with $P \otimes B^{\otimes l}$. Knowing the maps in question, we can deduce the following result essentially by direct calculation.

PROPOSITION 7.2. *Assuming the identity $H^0(H_n, P \otimes B^{\otimes l}) = R(n,l)$ and the vanishing of higher cohomology for these sheaves, we have*

$$H^0(Z_n, P \otimes B^{\otimes l}) \cong R(n,l)/M,$$

$$H^i(Z_n, P \otimes B^{\otimes l}) = 0 \quad \textit{for } i > 0,$$

where $M = R(n,l)\mathfrak{m}$ is the ideal generated by the homogeneous maximal ideal $\mathfrak{m} = \mathbb{C}[\mathbf{x},\mathbf{y}]_+^{S_n}$ in the subring $\mathbb{C}[\mathbf{x},\mathbf{y}]^{S_n} \subseteq R(n,l)$.

Note that the elements of $\mathfrak{m}$ represent global functions on H_n that vanish on Z_n, so the ideal M is necessarily contained in the kernel of the canonical map $R(n,l) \to H^0(Z_n, P \otimes B^{\otimes l})$ given by restriction of sections from H_n to Z_n. The content of the proposition is that this canonical map is surjective and its kernel is exactly M. The resolution (22) can be combined with the Atiyah–Bott formula for each of its terms to get the character formula for $R(n,l)/M$, as follows.

$$\mathcal{F}_{R(n,l)/M}(z;q,t) = \sum_{|\mu|=n} \frac{(1-q)(1-t)\Pi_\mu(q,t)B_\mu(q,t)^{l+1}\tilde{H}_\mu(z)}{\prod_{x\in\mu}(1-t^{-l(x)}q^{1+a(x)})(1-t^{1+l(x)}q^{-a(x)})}, \quad (23)$$

where

$$\Pi_\mu(q,t) = \prod_{(r,s)\in\mu\setminus\{(0,0)\}} (1-t^r q^s) = e_{n-1}[B_\mu - 1].$$

This is particularly interesting in the case $l = 0$, where $R(n,l)/M = \mathbb{C}[\mathbf{x},\mathbf{y}]/M$, and $M = (\mathbb{C}[\mathbf{x},\mathbf{y}]_+^{S_n})$, so $R(n,l)/M$ is the ring of coinvariants for the diagonal action of S_n on $\mathbb{C}^{2n}$. Its Frobenius series is equal to that of the diagonal harmonics DH_n, so we have shown that $\mathcal{F}_{DH_n}(z;q,t)$ is given by formula (23) with $l=0$. In [10], Garsia and I showed that this reduces to

$$\mathcal{F}_{DH_n}(z;q,t) = \nabla e_n,$$

where ∇ is the operator given in each degree n by $\nabla_{(1^n)}$ in our previous notation, or simply

$$\nabla \tilde{H}_\mu(z) = t^{n(\mu)} q^{n(\mu')} \tilde{H}_\mu.$$

This result follows from the expansion of e_n in terms of Macdonald polynomials, which can be calculated using the Cauchy formula for Macdonald polynomials and other facts from the theory of symmetric functions. Using the known specializations of $\tilde{H}_\mu(z;q,t)$ for $q=1$ and for $q=t^{-1}$, we were also able to calculate explicitly

$$\mathcal{F}_{DH_n}(z;1,t) \quad \text{and} \quad \mathcal{F}_{DH_n}(z;q,q^{-1}),$$

thereby establishing a whole series of combinatorial conjectures by ourselves and others (recounted in [14]), relating the character of DH_n to parking functions, trees, q-Lagrange inversion and the like. In particular, using either specialization to set both q and t equal to 1, we obtain Theorem 1.4, as stated in the first lecture:

$$\dim DH_n = (n+1)^{n-1}.$$

In our paper, of course, we proved all this modulo the *master conjecture* that $\mathcal{F}_{DH_n}(z;q,t) = \nabla e_n$. Once we justify the assumptions referred to in Proposition 7.2, the master conjecture follows.

7.3. VARIATIONS ON THE ZERO-FIBER

Just as we did for $R(n,l)$, we can enquire for the description of $R(n,l)/M$ as an $(S_n \times S_l)$-module. As before, we define

$$(R(n,l)/M)^\nu = \operatorname{Hom}_{S_l}(V^\nu, R(n,l)/M).$$

We conclude as before that its character formula is given by (23), with the factor B_μ^{l+1} in the numerator replaced by $B_\mu S_\nu[B_\mu]$. This yields ∇_ν applied to the case $l=1$ of (23), or

$$\mathcal{F}_{(R(n,l)/M)^\nu} = \nabla_\nu \nabla e_n.$$

In particular, this expression is a q,t-Schur positive symmetric function with polynomial coefficients (since $R(n,l)/M$ is finite dimensional). In [1], Garsia, F. Bergeron, Tesler and I conjectured, among other things, the stronger assertion that

$$\nabla_\nu e_n$$

is q,t-Schur positive for all ν. Operating on the Atiyah–Bott formula with ∇^{-1} multiplies the numerator by $t^{-n(\mu)}q^{-n(\mu')} = 1/e_n[B_\mu]$, which is the $\mathbb{T}^2$-character of the fiber of the line bundle $\mathcal{O}(-1) = \wedge^n B^*$. Hence the q,t-Schur positivity of $\nabla_\nu e_n$ would be a consequence of the following improvement of our vanishing theorem.

CONJECTURE 7.3. *We have*

$$H^i(H_n, \mathcal{O}(-1) \otimes P \otimes B^{\otimes l}) = 0 \quad \textit{for } i > 0.$$

Note that the perfect multiplication pairing $P \otimes P \to \mathcal{O}(1)$ implies that $\mathcal{O}(-1) \otimes P$ is the dual bundle P^*, so we can also state the conjecture as

$$H^i(H_n, P^* \otimes B^{\otimes l}) = 0 \quad \text{for } i > 0.$$

It may be possible to prove this by similar techniques to the vanishing theorem for $P \otimes B^{\otimes l}$, but I haven't succeeded in doing so yet. I also don't know for sure what space should be identified with $H^0(H_n, P^* \otimes B^{\otimes l})$.

7.4. VANISHING THEOREM

To finish, I want to explain how to prove the vanishing theorem we have been relying upon.

THEOREM 7.4. *The canonical map*

$$R(n, l) \to H^0(H_n, P \otimes B^{\otimes l})$$

is an isomorphism, and we have

$$H^i(H_n, P \otimes B^{\otimes l}) \quad \textit{for } i > 0.$$

This theorem is a consequence of a remarkable recent result of Bridgeland, King and Reid [3], combined with the results in [17]. The Bridgeland–King–Reid theorem concerns the *Hilbert scheme of orbits* $V /\!\!/ G$, for a finite group G acting on a complex vector space V as a subgroup of $\mathrm{SL}(V)$. For all vectors v in a Zariski open subset of V, the stabilizer of v is trivial. The orbit Gv then has $|G|$ distinct points and represents a point of the Hilbert scheme $\mathrm{Hilb}^{|G|}(V)$. By definition, $V /\!\!/ G$ is the closure in $\mathrm{Hilb}^G(V)$ of the locus formed by such points. There is a canonical "Chow morphism"

$$\sigma \colon V /\!\!/ G \to V/G,$$

and a commutative diagram

$$\begin{array}{ccc} X & \xrightarrow{f} & V \\ {\scriptstyle\rho}\downarrow & & \downarrow \\ V /\!\!/ G & \xrightarrow{\sigma} & V/G, \end{array}$$

where X is the universal family over $V /\!\!/ G$. By definition, $V /\!\!/ G$ is reduced and irreducible, and X is flat over $V /\!\!/ G$, so X is also reduced. Thus X is the reduced fiber product in the above diagram. Let $D(V /\!\!/ G)$ be the derived category of complexes of coherent sheaves on $V /\!\!/ G$ with bounded, coherent cohomology, and let $D^G(V)$ be the bounded, coherent derived category of G-equivariant complexes of sheaves on V. Equivalently, if $\mathbb{C}[\mathbf{z}]$ denotes the coordinate ring of V, we can identify $D^G(V)$ with the derived category of bounded complexes of finitely-generated G-equivariant $\mathbb{C}[\mathbf{z}]$-modules. Consider the derived functor

$$\Phi = Rf_* \circ \rho^* \colon D(V /\!\!/ G) \to D^G(V).$$

The result of Bridgeland–King–Reid is the following remarkable theorem.

THEOREM 7.5. *Assuming a certain condition on the dimensions of the fibers of $\sigma \colon V /\!\!/ G \to V/G$ holds, we have*

(1) σ is a crepant resolution of singularities, that is, $V /\!\!/ G$ is non-singular and has trivial canonical sheaf, and
(2) the functor $\Phi \colon D(V /\!\!/ G) \to D^G(V)$ is an equivalence of categories.

We apply this with $V = \mathbb{C}^{2n}$, $G = S_n$. In this case we know that X_n is a flat family over H_n whose generic fibers describe regular S_n-orbits in $\mathbb{C}^{2n}$. It follows from the universal property of $V /\!\!/ G$ that we get a map

$$H_n \to \mathbb{C}^{2n} /\!\!/ S_n$$

commuting with the Chow morphisms. It is elementary to construct the inverse map, so we have an isomorphism

$$\mathbb{C}^{2n} /\!\!/ S_n \cong H_n$$

identifying X_n with the universal family over $\mathbb{C}^{2n} /\!\!/ S_n$. This shows directly that $\mathbb{C}^{2n} /\!\!/ S_n$ is a crepant resolution of $\mathbb{C}^{2n}/S_n = S^n\mathbb{C}^2$, and the Bridgeland–King–Reid fiber dimension condition is easy to verify for H_n. Hence the functor

$$\Phi \colon D(H_n) \to D^{S_n}(\mathbb{C}^{2n})$$

is an equivalence of categories. Identifying $D^{S_n}(\mathbb{C}^{2n})$ with the derived category of S_n-equivariant $\mathbb{C}[\mathbf{x}, \mathbf{y}]$-modules, the functor Φ becomes

$$\Phi = R\Gamma(P \otimes -).$$

In this terminology, since the cohomology functors $H^i(H_n, -)$ are the derived functors $R^i\Gamma$, Theorem 7.4 reduces to the identity

$$\Phi(B^{\otimes l}) = R(n, l).$$

To prove this, we can equally well prove

$$\Psi R(n, l) = B^{\otimes l},$$

where Ψ is the inverse functor. In the proof of their theorem, Bridgeland–King–Reid give Ψ explicitly as

$$\Psi = (\rho_*(\omega_X \underset{L}{\otimes} Lf^* -))^G.$$

In our situation ω_{X_n} is the line bundle $\mathcal{O}(-1)$ pulled back from H_n, and the S_n action on it is the sign character tensored with the canonical trivial action. For us, Ψ is therefore given by

$$\Psi = (\rho_*(\mathcal{O}(-1) \otimes Lf^* -))^\epsilon.$$

We have a canonical map $R(n,l) \to \Phi(B^{\otimes l})$ and a corresponding map $\Psi R(n,l) \to B^{\otimes l}$, and we are to show that the latter is an isomorphism in the derived category. In the derived category there is a distinguished triangle on the map $\Psi R(n,l) \to B^{\otimes l}$, and we are equivalently to show that its apex C is zero. It follows from the freeness theorem for $R(n,l)$, Theorem 1.8, that $R(n,l)$ has an S_n-equivariant $\mathbb{C}[\mathbf{x}, \mathbf{y}]$-free resolution of length $n-1$. Using this resolution to compute $Lf^* R(n,l)$, we find that $\Psi R(n,l)$ has a locally free resolution of length $n-1$, and hence C has a locally free resolution of length n. In other words, C is represented as an object in the derived category by a complex of length n of locally free sheaves, and we are to show that this complex is exact.

A fundamental result of commutative algebra, known as the *new intersection theorem*, implies that the cohomology of the complex C, if not zero, has support of codimension at most n. Outside the locus $Z_n \times \mathbb{C}^2$ of points $I \in H_n$ such that $\sigma(I)$ is concentrated at a single point P, we can use induction on n to show that C is locally exact. Thus the support of C is contained in the $(n+1)$-dimensional locus $Z_n \times \mathbb{C}^2$. To show that C is exact, we need to cut its support down by two more dimensions. We do this by an explicit calculation showing that C is locally exact on the curvilinear locus. This suffices, because the non-curvilinear locus has codimension two in Z_n, and completes the proof of Theorem 7.4.

References

1. F. Bergeron, A. M. Garsia, M. Haiman, and G. Tesler, *Identities and positivity conjectures for some remarkable operators in the theory of symmetric functions*, Methods Appl. Anal. **6** (1999), no. 3, 363–420.
2. N. Bergeron and A. M. Garsia, *On certain spaces of harmonic polynomials*, Hypergeometric functions on domains of positivity, Jack polynomials, and applications (Tampa, FL, 1991), Amer. Math. Soc., Providence, RI, 1992, pp. 51–86.
3. Tom Bridgeland, Alastair King, and Miles Reid, *The McKay correspondence as an equivalence of derived categories*, J. Amer. Math. Soc. **14** (2001), 535–554.
4. Jan Cheah, *Cellular decompositions for nested Hilbert schemes of points*, Pacific J. Math. **183** (1998), no. 1, 39–90.

5. Geir Ellingsrud and Stein Arild Strømme, *On the homology of the Hilbert scheme of points in the plane*, Invent. Math. **87** (1987), no. 2, 343–352.
6. Jacques Emsalem, *Géométrie des points épais*, Bull. Soc. Math. France **106** (1978), no. 4, 399–416.
7. John Fogarty, *Algebraic families on an algebraic surface*, Amer. J. Math **90** (1968), 511–521.
8. A. M. Garsia and J. Haglund, *A positivity result in the theory of Macdonald polynomials*, Proc. Natl. Acad. Sci. USA **98** (2001), no. 8, 4313–4316 (electronic).
9. ———, *A proof of the q,t-Catalan positivity conjecture*, Discrete Mathematics (to appear), Proceedings of the September 2000 Montreal conference on Algebraic Combinatorics.
10. A. M. Garsia and M. Haiman, *A remarkable q,t-Catalan sequence and q-Lagrange inversion*, J. Algebraic Combin. **5** (1996), no. 3, 191–244.
11. A. M. Garsia and C. Procesi, *On certain graded S_n-modules and the q-Kostka polynomials*, Adv. Math. **94** (1992), no. 1, 82–138.
12. A. M. Garsia and J. Remmel, *Plethystic formulas and positivity for q,t-Kostka coefficients*, Mathematical essays in honor of Gian-Carlo Rota (Cambridge, MA, 1996), Birkhäuser Boston, Boston, MA, 1998, pp. 245–262.
13. A. M. Garsia and G. Tesler, *Plethystic formulas for Macdonald q,t-Kostka coefficients*, Adv. Math. **123** (1996), no. 2, 144–222.
14. Mark Haiman, *Conjectures on the quotient ring by diagonal invariants*, J. Algebraic Combin. **3** (1994), no. 1, 17–76.
15. ———, *t,q-Catalan numbers and the Hilbert scheme*, Discrete Math. **193** (1998), no. 1-3, 201–224, Selected papers in honor of Adriano Garsia (Taormina, 1994).
16. ———, *Macdonald polynomials and geometry*, New perspectives in geometric combinatorics (Billera, Björner, Greene, Simion, and Stanley, eds.), MSRI Publications, vol. 38, Cambridge University Press, 1999, pp. 207–254.
17. ———, *Hilbert schemes, polygraphs, and the Macdonald positivity conjecture.*, J. Amer. Math. Soc. **14** (2001), 941–1006.
18. ———, *Vanishing theorems and character formulas for the Hilbert scheme of points in the plane*, Preprint, UCB (2001).
19. ———, *Vanishing theorems and character formulas for the Hilbert scheme of points in the plane (abbreviated version)*, Physics and Combinatorics 2000 (A. N. Kirillov and Nadejda Liskova, eds.), World Scientific, 2001, Proceedings of the Nagoya 2000 International Workshop, pp. 1–21.
20. R. Hotta and T. A. Springer, *A specialization theorem for certain Weyl group representations and an application to the Green polynomials of unitary groups*, Invent. Math. **41** (1977), no. 2, 113–127.
21. Anatol N. Kirillov and Masatoshi Noumi, *Affine Hecke algebras and raising operators for Macdonald polynomials*, Duke Math. J. **93** (1998), no. 1, 1–39.
22. Friedrich Knop, *Integrality of two variable Kostka functions*, J. Reine Angew. Math. **482** (1997), 177–189.
23. Alain Lascoux and Marcel-Paul Schützenberger, *Sur une conjecture de H. O. Foulkes*, C. R. Acad. Sci. Paris Sér. A-B **286** (1978), no. 7, A323–A324.
24. I. G. Macdonald, *Symmetric functions and Hall polynomials*, second ed., The Clarendon Press, Oxford University Press, New York, 1995, With contributions by A. Zelevinsky, Oxford Science Publications.
25. Siddhartha Sahi, *Interpolation, integrality, and a generalization of Macdonald's polynomials*, Internat. Math. Res. Notices **10** (1996), 457–471.

26. T. A. Springer, *A construction of representations of Weyl groups*, Invent. Math. **44** (1978), no. 3, 279–293.
27. R. W. Thomason, *Une formule de Lefschetz en K-théorie équivariante algébrique*, Duke Math. J. **68** (1992), no. 3, 447–462.