FOUNDATIONS OF ALGEBRAIC GEOMETRY CLASS 31

RAVI VAKIL

CONTENTS

Last day: Cohomology continued. Beginning of Hilbert functions

Today: Hilbert polynomials and Hilbert functions. Higher direct image sheaves.

1. APPLICATION OF COHOMOLOGY: HILBERT POLYNOMIALS AND HILBERT FUNCTIONS; DEGREES

We're in the process of seeing applications of cohomology. In this section, we will work over a field k. We defined $h^i(X, \mathcal{F}) := \dim_k H^i(X, \mathcal{F})$. If $\mathcal{F}$ is a coherent sheaf on a projective k-scheme X, we defined the *Euler characteristic*

$$\chi(X, \mathcal{F}) = \sum_{i=0}^{\dim X} (-1)^i h^i(X, \mathcal{F}).$$

We will see repeatedly here and later that Euler characteristics behave better than individual cohomology groups.

If $\mathcal{F}$ is a coherent sheaf on X, define the *Hilbert function of* $\mathcal{F}$:

$$h_{\mathcal{F}}(m) := h^0(X, \mathcal{F}(m)).$$

The *Hilbert function of* X is the Hilbert function of the structure sheaf $\mathcal{O}_X$. The ancients were aware that the Hilbert function is "eventually polynomial", i.e. for large enough n, it agrees with some polynomial, called the *Hilbert polynomial* (and denoted $p_{\mathcal{F}}(m)$ or $p_X(m)$). In modern language, we expect that this is because the Euler characteristic should be a polynomial, and that for $m \gg 0$, the higher cohomology vanishes. This is indeed the case, as we now verify.

Date: Tuesday, February 14, 2006. Tiny update April 9.

1.1. *Theorem. — If $\mathcal{F}$ is a coherent sheaf on a projective k-scheme $X \hookrightarrow \mathbb{P}^n_k$, for $m \gg 0$, $h^0(X, \mathcal{F}(m))$ is a polynomial of degree equal to the dimension of the support of $\mathcal{F}$. In particular, for $m \gg 0$, $h^0(X, \mathcal{O}_X(m))$ is polynomial with degree* = $\dim X$.

(Here $\mathcal{O}_X(m)$ is the restriction or pullback of $\mathcal{O}_{\mathbb{P}^n_k}(1)$.)

I realize now that I will use the notion of associated primes of a *module*, not just of a ring. I think I only discussed associated primes of a ring last quarter, because I had hoped not to need this slightly more general case. Now I really don't need it, and if you want to ignore this issue, you can just prove the second half of the theorem, which is all we will use anyway. But the argument carries through with no change, so please follow along if you can.

Proof. For $m \gg 0$, $h^i(X, \mathcal{F}(m)) = 0$ by Serre vanishing (class 29 Theorem 4.2(ii)), so instead we will prove that for *all* m, $\chi(X, \mathcal{F}(m))$ is a polynomial of degree equal to the dimension of the support of $\mathcal{F}$. Define $p_{\mathcal{F}}(m) = \chi(X, \mathcal{F}(m))$; we'll show that $p_{\mathcal{F}}(m)$ is a polynomial of the desired degree.

Our approach will be a little weird. We'll have two steps, and they will be very similar. If you can streamline, please let me know.

Step 1. We first show that for all n, if $\mathcal{F}$ is scheme-theoretically supported a linear subspace of dimension k (i.e. $\mathcal{F}$ is the pushforward of a coherent sheaf on some linear subspace of dimension k), then $p_{\mathcal{F}}(m)$ is a polynomial of degree at most k. (In particular, for any coherent $\mathcal{F}$, $p_{\mathcal{F}}(m)$ is a polynomial of degree at most n.)

We prove this by induction on the dimension of the support. I'll leave the base case $k = 0$ (or better yet, $k = -1$) to you (*exercise*). Suppose now that X is supported in a linear space Λ of dimension k, and we know the result for all $k' < k$. Then let $x = 0$ be a hyperplane not containing Λ, so $\Lambda' = \dim(x = 0) \cap \Lambda = k - 1$. Then we have an exact sequence

$$0 \longrightarrow \mathcal{K} \longrightarrow \mathcal{F} \xrightarrow{\times x} \mathcal{F}(1) \longrightarrow \mathcal{K}' \longrightarrow 0 \qquad (1)$$

where $\mathcal{K}$ (resp. $\mathcal{K}'$) is the kernel (resp. cokernel) of the map $\times x$. Notice that $\mathcal{K}$ and $\mathcal{K}'$ are both supported on Λ'. (This corresponds to an algebraic fact: over an affine open $\operatorname{Spec} A$, the exact sequence is

$$0 \longrightarrow K \longrightarrow M \xrightarrow{\times x} M \longrightarrow K' \longrightarrow 0$$

and both $K = \ker(\times x) = (0 : x)$ and $K' \cong M/xM$ are (A/x)-modules.) Twist (1) by $\mathcal{O}(m)$ and take Euler-characteristics to obtain $p_{\mathcal{F}}(m+1) - p_{\mathcal{F}}(m) = p_{\mathcal{K}'}(m) - p_{\mathcal{K}}(m)$. By the inductive hypothesis, the right side of this equation is a polynomial of degree at most $k - 1$. Hence (by an easy induction) $p(m)$ is a polynomial of degree at most k.

Step 2. We'll now show that the degree of this polynomial is precisely $\dim \operatorname{Supp} \mathcal{F}$. As $\mathcal{F}$ is a coherent sheaf on a Noetherian scheme, it has a finite number of associated points, so we can find a hypersurface $H = (f = 0)$ not containing any of the associated points. (This is that problem from last quarter that we have been repeatedly using recently: problem 24(c) on set 5, which was exercise 1.19 in the class 11 notes.) In particular, $\dim H \cap \operatorname{Supp} \mathcal{F}$

is strictly less than $\dim \operatorname{Supp} \mathcal{F}$, and in fact one less by Krull's Principal Ideal Theorem. Let $d = \deg f$. Then I claim that $\times f : \mathcal{F}(-d) \to \mathcal{F}$ is an inclusion. Indeed, on any affine open set, the map is of the form $\times \overline{f} : M \to M$ (where $\overline{f}$ is the restriction of f to this open set), and the fact that $f = 0$ contains no associated points *means* that this is an injection of modules. (Remember that those ring elements annihilating elements of M are precisely the associated primes, and $\overline{f}$ is contained in none of them.) Then we have

$$0 \to \mathcal{F}(-d) \to \mathcal{F} \to \mathcal{K}' \to 0.$$

Twisting by $\mathcal{O}(m)$ yields

$$0 \to \mathcal{F}(m-d) \to \mathcal{F}(m) \to \mathcal{K}'(m) \to 0.$$

Taking Euler characteristics gives $p_{\mathcal{F}}(m) - p_{\mathcal{F}}(m-d) = p_{\mathcal{K}'}(m)$. Now by step 1, we know that $p_{\mathcal{F}}(m)$ is a polynomial. Also, by our inductive hypothesis, and Exercise 1.2 below, the right side is a polynomial of degree of precisely $\dim \operatorname{Supp} \mathcal{F} - 1$. Hence $p(m)$ is a polynomial of degree $\dim \operatorname{Supp} \mathcal{F}$. □

1.2. *Exercise.* Consider the short exact sequence of A-modules $0 \longrightarrow M \xrightarrow{\times f} M \longrightarrow K' \longrightarrow 0$. Show that $\operatorname{Supp} K' = \operatorname{Supp}(M) \cap \operatorname{Supp}(A/f)$.

Notice that we needed the first part of the proof to ensure that $p_{\mathcal{F}}(m)$ is in fact a polynomial; otherwise, the second part would just show that $p_{\mathcal{F}}(m)$ is just a polynomial when m is fixed modulo d.

(For experts: here is a different way to avoid having two similar steps. If k is an infinite field, e.g. if it were algebraically closed, then we could find a hypersurface as in step 2 of degree 1, using that problem from last quarter mentioned in the proof. So what to do if k is not infinite? Note that if you have a complex of k-vector spaces, and you take its cohomology, and then tensor with $\overline{k}$, you get the same thing as if you tensor first, and then take the cohomology. By this trick, we can assume that k is algebraically closed. In fancy language: we have taken a *faithfully flat* base extension. I won't define what this means here; it will turn up early in the third quarter.)

Example 1. $p_{\mathbb{P}^n}(m) = \binom{m+n}{n}$, where we interpret this as the polynomial $(m+1)\cdots(m+n)/n!$.

Example 2. Suppose H is a degree d hypersurface in $\mathbb{P}^n$. Then from

$$0 \to \mathcal{O}_{\mathbb{P}^n}(-d) \to \mathcal{O}_{\mathbb{P}^n} \to \mathcal{O}_H \to 0,$$

we have

$$p_H(m) = p_{\mathbb{P}^n}(m) - p_{\mathbb{P}^n}(m-d) = \binom{m+n}{n} - \binom{m+n-d}{n}.$$

1.3. *Exercise.* Show that the twisted cubic (in $\mathbb{P}^3$) has Hilbert polynomial $3m+1$.

1.4. *Exercise.* Find the Hilbert polynomial for the dth Veronese embedding of $\mathbb{P}^n$ (i.e. the closed immersion of $\mathbb{P}^n$ in a bigger projective space by way of the line bundle $\mathcal{O}(d)$).

From the Hilbert polynomial, we can extract many invariants, of which two are particularly important. The first is the *degree*. Classically, the degree of a complex projective variety of dimension n was defined as follows. We slice the variety with n generally chosen hyperplane. Then the intersection will be a finite number of points. The degree is this number of points. Of course, this requires showing all sorts of things. Instead, we will define the *degree of a projective k-scheme of dimension* n to be leading coefficient of the Hilbert polynomial (the coefficient of m^n) times $n!$.

For example, the degree of $\mathbb{P}^n$ in itself is 1. The degree of the twisted cubic is 3.

1.5. *Exercise.* Show that the degree of a degree d hypersurface is d (preventing a notational crisis).

1.6. *Exercise.* Suppose a curve C is embedded in projective space via an invertible sheaf of degree d. (In other words, this line bundle determines a closed immersion.) Show that the degree of C under this embedding is d (preventing another notational crisis). (Hint: Riemann-Roch.)

1.7. *Exercise.* Find the degree of the dth Veronese embedding of $\mathbb{P}^n$.

1.8. *Exercise (Bezout's theorem).* Suppose X is a projective scheme of dimension at least 1, and H is a degree d hypersurface not containing any associated points of X. (For example, if X is a projective variety, then we are just requiring H not to contain any irreducible components of X.) Show that $\deg H \cap X = d \deg X$.

This is a very handy theorem! For example: if two projective plane curves of degree m and degree n share no irreducible components, then they intersect in mn points, counted with appropriate multiplicity. The notion of multiplicity of intersection is just the degree of the intersection as a k-scheme.

We trot out a useful example for a third time: let $k = \mathbb{Q}$, and consider the parabola $x = y^2$. We intersect it with the four usual suspects: $x = 1$, $x = 0$, $x = -1$, and $x = 2$, and see that we get 2 each time (counted with the same convention as with the last time we saw this example).

If we intersect it with $y = 2$, we only get one point — but that's of course because this isn't a projective curve, and we really should be doing this intersection on $\mathbb{P}^2_k$ — and in this case, the conic meets the line in two points, one of which is "at ∞".

1.9. *Exercise.* Determine the degree of the d-fold Veronese embedding of $\mathbb{P}^n$ in a different way as follows. Let $\nu_d : \mathbb{P}^n \to \mathbb{P}^N$ be the Veronese embedding. To find the degree of the image, we intersect it with n hyperplanes in $\mathbb{P}^N$ (scheme-theoretically), and find the number of intersection points (counted with multiplicity). But the pullback of a hyperplane in $\mathbb{P}^N$ to $\mathbb{P}^n$ is a degree d hypersurface. Perform this intersection in $\mathbb{P}^n$, and use Bezout's theorem. (If already you know the answer by the earlier exercise on the degree of the Veronese embedding, this will be easier.)

There is another nice bit of information residing in the Hilbert polynomial. Notice that $p_X(0) = \chi(X, \mathcal{O}_X)$, which is an *intrinsic* invariant of the scheme X, which does not depend on the projective embedding.

Imagine how amazing this must have seemed to the ancients: they defined the Hilbert function by counting how many "functions of various degrees" there are; then they noticed that when the degree gets large, it agrees with a polynomial; and then when they plugged 0 into the polynomial — extrapolating backwards, to where the Hilbert function and Hilbert polynomials didn't agree — they found a magic invariant!

And now I can give you a nonsingular curve over an algebraically closed field that is not $\mathbb{P}^1$! Note that the Hilbert polynomial of $\mathbb{P}^1$ is $(m+1)/1 = m+1$, so $\chi(\mathcal{O}_{\mathbb{P}^1}) = 1$. Suppose C is a degree d curve in $\mathbb{P}^2$. Then the Hilbert polynomial of C is

$$p_{\mathbb{P}^2}(m) - p_{\mathbb{P}^2}(m-d) = (m+1)(m+2)/2 - (m-d+1)(m-d+2)/2.$$

Plugging in $m = 0$ gives us $-(d^2 - 3d)/2$. Thus when $d > 2$, we have a curve that cannot be isomorphic to $\mathbb{P}^1$! (I think I gave you an earlier exercise that there is a *nonsingular* degree d curve. Note however that the calculation above didn't use nonsingularity.)

Now from $0 \to \mathcal{O}_{\mathbb{P}^2}(-d) \to \mathcal{O}_{\mathbb{P}^2} \to \mathcal{O}_C \to 0$, using $h^1(\mathcal{O}_{\mathbb{P}^2}(d)) = 0$, we have that $h^0(C, \mathcal{O}_C) = 1$. As $h^0 - h^1 = \chi$, we have

$$h^1(C, \mathcal{O}_C) = (d-1)(d-2)/2.$$

Motivated by geometry, we define the *arithmetic genus* of a scheme X as $1 - \chi(X, \mathcal{O}_X)$. This is sometimes denoted $p_a(X)$. In the case of nonsingular complex curves, this corresponds to the topological genus. For irreducible reduced curves (or more generally, curves with $h^0(X, \mathcal{O}_X) \cong k$), $p_a(X) = h^1(X, \mathcal{O}_X)$. (In higher dimension, this is a less natural notion.)

We thus now have examples of curves of genus $0, 1, 3, 6, 10, \ldots$ (corresponding to degree 1 or $2, 3, 4, 5, \ldots$).

This begs some questions, such as: are there curves of other genera? (Yes.) Are there other genus 1 curves? (Not if k is algebraically closed, but yes otherwise.) Do we have all the curves of genus 3? (Almost all, but not quite all.) Do we have all the curves of genus 6? (We're missing most of them.)

Caution: The Euler characteristic of the structure sheaf doesn't distinguish between isomorphism classes of nonsingular projective schemes over algebraically closed fields — for example, $\mathbb{P}^1 \times \mathbb{P}^1$ and $\mathbb{P}^2$ both have Euler characteristic 1, but are not isomorphic (as for example $\operatorname{Pic} \mathbb{P}^2 \cong \mathbb{Z}$ while $\operatorname{Pic} \mathbb{P}^1 \times \mathbb{P}^1 \cong \mathbb{Z} \oplus \mathbb{Z}$).

Important Remark. We can restate the Riemann-Roch formula as:

$$h^0(C, \mathcal{L}) - h^1(C, \mathcal{L}) = \deg \mathcal{L} - p_a + 1.$$

This is the most common formulation of the Riemann-Roch formula.

1.10. *Complete intersections.* We define a *complete intersection* in $\mathbb{P}^n$ as follows. $\mathbb{P}^n$ is a complete intersection in itself. A closed subscheme $X_r \hookrightarrow \mathbb{P}^n$ of dimension r (with $r < n$) is a complete intersection if there is a complete intersection X_{r+1}, and X_r is a Cartier divisor in class $\mathcal{O}_{X_{r+1}}(d)$.

Exercise. Show that if X is a complete intersection of dimension r in $\mathbb{P}^n$, then $H^i(X, \mathcal{O}_X(m)) = 0$ for all $0 < i < r$ and all m. Show that if $r > 0$, then $H^0(\mathbb{P}^n, \mathcal{O}(m)) \to H^0(X, \mathcal{O}(m))$ is surjective.

Now in my definition, X_r is the zero-divisor of a section of $\mathcal{O}_{X_{r+1}}(m)$ for some m. But this section is the restriction of a section of $\mathcal{O}(m)$ on $\mathbb{P}^n$. Hence X_r is the scheme-theoretic intersection of X_{r+1} with a hypersurface. Thus inductively we can show that X_r is the scheme-theoretic intersection of $n - r$ hypersurfaces. (By Bezout's theorem, $\deg X_r$ is the product of the degree of the defining hypersurfaces.)

Exercise. Show that complete intersections of positive dimension are connected. (Hint: show $h^0(X, \mathcal{O}_X) = 1$.)

Exercise. Find the genus of the intersection of 2 quadrics in $\mathbb{P}^3$. (We get curves of more genera by generalizing this!)

Exercise. Show that the rational normal curve of degree d in $\mathbb{P}^d$ is *not* a complete intersection if $d > 2$.

Exercise. Show that the union of 2 distinct planes in $\mathbb{P}^4$ is not a complete intersection. (This is the first appearance of another universal counterexample!) Hint: it is connected, but you can slice with another plane and get something not connected.

This is another important scheme in algebraic geometry that is an example of many sorts of behavior. We will see more of it later!

2. Higher direct image sheaves

I'll now introduce a notion generalizing these Cech cohomology groups. Cohomology groups were defined for $X \to \operatorname{Spec} A$ where the structure morphism is quasicompact and separated; for any quasicoherent $\mathcal{F}$ on X, we defined $H^i(X, \mathcal{F})$.

We'll now do something similar for quasicompact and separated morphisms $\pi : X \to Y$: for any quasicoherent $\mathcal{F}$ on X, we'll define $R^i\pi_*\mathcal{F}$, a quasicoherent sheaf on Y.

We have many motivations for doing this. In no particular order:

(1) It "globalizes" what we were doing anywhere.
(2) If $0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0$ is a short exact sequence of quasicoherent sheaves on X, then we know that $0 \to \pi_*\mathcal{F} \to \pi_*\mathcal{G} \to \pi_*\mathcal{H}$ is exact, and higher pushforwards will extend this to a long exact sequence.

(3) We'll later see that this will show how cohomology groups vary in families, especially in "nice" situations. Intuitively, if we have a nice family of varieties, and a family of sheaves on them, we could hope that the cohomology varies nicely in families, and in fact in "nice" situations, this is true. (As always, "nice" usually means "flat", whatever that means.)

There will be no extra work involved for us.

Suppose $\pi : X \to Y$, and $\mathcal{F}$ is a quasicoherent sheaf on X. For each $\operatorname{Spec} A \subset Y$, we have A-modules $H^i(\pi^{-1}(\operatorname{Spec} A), \mathcal{F})$. We will show that these patch together to form a quasicoherent sheaf. We need check only one fact: that this behaves well with respect to taking distinguished opens. In other words, we must check that for each $f \in A$, the natural map $H^i(\pi^{-1}(\operatorname{Spec} A), \mathcal{F}) \to H^i(\pi^{-1}(\operatorname{Spec} A), \mathcal{F})_f$ (induced by the map of spaces in the opposite direction — H^i is contravariant in the space) is precisely the localization $\otimes_A A_f$. But this can be verified easily: let $\{U_i\}$ be an affine cover of $\pi^{-1}(\operatorname{Spec} A)$. We can compute $H^i(\pi^{-1}(\operatorname{Spec} A), \mathcal{F})$ using the Cech complex. But this induces a cover $\operatorname{Spec} A_f$ in a natural way: If $U_i = \operatorname{Spec} A_i$ is an affine open for $\operatorname{Spec} A$, we define $U_i' = \operatorname{Spec}(A_i)_f$. The resulting Cech complex for $\operatorname{Spec} A_f$ is the localization of the Cech complex for $\operatorname{Spec} A$. As taking cohomology of a complex commutes with localization, we have defined a quasicoherent sheaf on Y by one of our definitions of quasicoherent sheaves.

2.1. (Something important happened in that last sentence — localization commuting with taking cohomology. If you want practice with this notion, here is an *exercise:* suppose $C^0 \to C^1 \to C^2$ is a complex in an abelian category, and F is an exact functor to another abelian category. Show that F applied to the cohomology of this complex is naturally isomorphic to the cohomology of F of this complex. Translation: taking cohomology commutes with exact functors. In the particular case of this construction, the exact functor in equation is the localization functor $\otimes_A A_f$ from A-modules to A_f-modules. I'll discuss this a bit more at the start of the class 32 notes.)

Define the i**th higher direct image sheaf** or the i**th (higher) pushforward sheaf** to be this quasicoherent sheaf.

2.2. *Theorem. —*

(a) $R^0\pi_*\mathcal{F}$ *is canonically isomorphic to* $\pi_*\mathcal{F}$.
(b) $R^i\pi_*$ *is a covariant functor from the category of quasicoherent sheaves on* X *to the category of quasicoherent sheaves on* Y*, and a contravariant functor in* Y*-schemes* X.
(c) *A short exact sequence* $0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0$ *of sheaves on* X *induces a long exact sequence*

$$0 \longrightarrow R^0\pi_*\mathcal{F} \longrightarrow R^0\pi_*\mathcal{G} \longrightarrow R^0\pi_*\mathcal{H} \longrightarrow$$

$$R^1\pi_*\mathcal{F} \longrightarrow R^1\pi_*\mathcal{G} \longrightarrow R^1\pi_*\mathcal{H} \longrightarrow \cdots$$

of sheaves on Y*. (This is often called the corresponding* **long exact sequence of higher pushforward sheaves***.)*

(d) *(projective pushforwards of coherent are coherent) If* π *is a projective morphism and* $\mathcal{O}_Y$ *is coherent on* Y *(this hypothesis is automatic for* Y *locally Noetherian), and* $\mathcal{F}$ *is a coherent sheaf on* X*, then for all* i, $R^i\pi_*\mathcal{F}$ *is a coherent sheaf on* Y.

Proof. Because it suffices to check each of these results on affine opens, they all follow from the analogous statements in Cech cohomology. □

The following result is handy (and essentially immediate from our definition).

2.3. *Exercise.* Show that if π is affine, then for $i > 0$, $R^i\pi_*\mathcal{F} = 0$. Moreover, show that if Y is quasicompact and quasiseparated then the natural morphism $H^i(X, \mathcal{F}) \to H^i(Y, f_*\mathcal{F})$ is an isomorphism. (A special case of the first sentence is a special case we showed earlier, when π is a closed immersion. Hint: use any affine cover on Y, which will induce an affine cover of X.)

E-mail address: vakil@math.stanford.edu