Lecture 19: April 17

Proof of Theorem 18.5. Today, we are going to finish the proof of Theorem 18.5. The statement is that, for any morphism $f\colon X \to Y$ between nonsingular algebraic varieties, one has:

(a) $f_+\colon D^b_h(\mathscr{D}_X^{op}) \to D^b_h(\mathscr{D}_Y^{op})$
(b) $\mathbf{L}f^*\colon D^b_h(\mathscr{D}_Y) \to D^b_h(\mathscr{D}_X)$

Last time, I sketched the argument that reduces both statements to the special case of a coordinate projection $p\colon \mathbb{A}^{n+1} \to \mathbb{A}^n$. Let me first fill in the proof of a crucial lemma that we used.

Lemma 19.1. *Let $i\colon X \hookrightarrow Y$ be a closed embedding of codimension one, and $j\colon U = Y \setminus X \hookrightarrow Y$ the complementary open embedding. Then for any holonomic right $\mathscr{D}_Y$-module $\mathcal{M}$, one has an exact sequence*

$$0 \to i_+(L^{-1}i^*\mathcal{M}) \to \mathcal{M} \to j_+(\mathcal{M}\big|_U) \to i_+(L^0 i^*\mathcal{M}) \to 0.$$

We had defined the pullback functor for left $\mathscr{D}$-modules. To compute $\mathbf{L}i^*\mathcal{M}$, one first converts $\mathcal{M}$ into a left $\mathscr{D}_Y$-module by $\mathcal{H}om_{\mathscr{D}_Y}(\mathscr{D}_Y^\omega, \mathcal{M})$, then applies the pullback functor $\mathbf{L}i^*$, and then converts the resulting left $\mathscr{D}_X$-module back into a right $\mathscr{D}_X$-module by tensoring with $\mathscr{D}_X^\omega$.

Proof. We are only going to prove the local version, since that is all that we need for the proof of Theorem 18.5. Suppose then that Y is affine, with coordinates $y_0, y_1, \dots, y_n$, and that X is the closed subscheme defined by $y_0 = 0$. Set $A = \Gamma(Y, \mathscr{O}_Y)$ and $M = \Gamma(Y, \mathcal{M})$, which is a holonomic right $D(A)$-module. After carrying out the left-right conversions, $\mathbf{L}i^*\mathcal{M}$ corresponds to the complex of $D(B)$-modules

$$M \xrightarrow{y_0} M \tag{19.2}$$

placed in degrees -1 and 0; here $B = \Gamma(X, \mathscr{O}_X)$. On the other hand, j is affine, and so $j_+(\mathcal{M}\big|_U) = j_*(\mathcal{M}\big|_U)$ is the localization

$$M \otimes_A A[y_0^{-1}].$$

We therefore have to analyze the kernel and cokernel of the natural morphism

$$\varphi\colon M \to M \otimes_A A[y_0^{-1}].$$

Let us first consider $\ker\varphi$. It consists of all $m \in M$ such that $my_0^\ell = 0$ for some $\ell \geq 1$. This submodule is supported on X, and by Kashiwara's equivalence, it is the direct image of a $D(B)$-module M_0. Here

$$M_0 = \big\{\, m \in M \;\big|\; my_0 = 0 \,\big\}$$

which is the $D(B)$-module corresponding to $L^{-1}i^*\mathcal{M}$ by (19.2). Next, we consider $\operatorname{coker}\varphi$. It consists of all finite sums of the form

$$\sum_{j \geq 0} m_j \otimes y_0^{-j},$$

with $m_j \in M$, modulo the image of M. This is again the direct image of a $D(B)$-module M_1, by Kashiwara's equivalence, where M_1 is the submodule annihilated by y_0. A short computation gives

$$M_1 = \big\{\, m_0 \otimes 1 + m_1 \otimes y_0^{-1} \;\big|\; m_0, m_1 \in M \,\big\}/M \cong M/My_0,$$

and again by (19.2), this is the $D(B)$-module corresponding to $L^0 i^*\mathcal{M}$. ∎

In fact, the lemma generalizes to arbitrary closed embeddings $i\colon X \hookrightarrow Y$. If we again let $j\colon U \hookrightarrow Y$ be the open embedding of the complement $U = Y \setminus X$, then we have a distinguished triangle (= short exact sequence)

$$\mathbf{R}\mathcal{H}_X(\mathscr{F}) \to \mathscr{F} \to \mathbf{R}j_*(\mathscr{F}\big|_U) \to \mathbf{R}\mathcal{H}_X(\mathscr{F})[1],$$

for every sheaf of $\mathscr{O}_Y$-modules $\mathscr{F}$, where where $\mathcal{H}_X$ is the functor of "sections with support in X". Concretely, $\mathbf{R}\mathcal{H}_X(\mathscr{F})$ is computed by choosing an injective resolution of $\mathscr{F}$ and applying the functor $\mathcal{H}_X$ to each sheaf in the resolution. When $\mathcal{M}$ is a right $\mathscr{D}_Y$-module, we have $\mathbf{R}j_*(\mathcal{M}\big|_U) = j_+(\mathcal{M}\big|_U)$, and the distinguished triangle becomes

$$\mathbf{R}\mathcal{H}_X(\mathcal{M}) \to \mathcal{M} \to j_+(\mathcal{M}\big|_U)\mathbf{R}\mathcal{H}_X(\mathcal{M})[1].$$

Then the fancy version of the lemma is that $\mathbf{R}\mathcal{H}_X(\mathcal{M})$ is isomorphic to $i_+\mathbf{R}i^*\mathcal{M}$, up to a shift by the codimension $\dim Y - \dim X$.

Coordinate projections. To prove Theorem 18.5, it remains to treat the case of a coordinate projection $p\colon \mathbb{A}^{n+1} \to \mathbb{A}^n$. We need to show that if $\mathcal{M}$ is a holonomic right $\mathscr{D}_{\mathbb{A}^{n+1}}$-module, then all cohomology sheaves of $p_+\mathcal{M}$ are holonomic $\mathscr{D}_{\mathbb{A}^n}$-modules. This brings us back to modules over the Weyl algebra. Let us first look at a concrete example.

Example 19.3. Consider the special case $p\colon \mathbb{A}^1 \to \operatorname{Spec} k$. The pushforward of a right A_1-module M is computed by the Spencer complex

$$M \xrightarrow{\partial} M$$

and the theorem is claiming that when M is holonomic, both the kernel and cokernel of multiplication by ∂ are finite-dimensional k-vector spaces. One approach would be to take a good filtration $F_\bullet M$ and pass to the associated graded $k[x,\partial]$-module $\operatorname{gr}^F M$. Its support is one-dimensional, but unfortunately, the kernel and cokernel of multiplication by ∂ can fail to be finite-dimensional. (This happens for example with $M = k[x]$.)

Let me show you an ad-hoc argument for why

$$\ker \partial = \left\{\, m \in M \;\middle|\; m\partial = 0 \,\right\}$$

has finite dimension over k. Consider the A_1-submodule

$$\ker \partial \cdot A_1 \subseteq M$$

generated by $\ker \partial$. Since M is finitely generated over A_1, this submodule is also finitely generated. The commutation relation $[\partial, x] = 1$ implies that, for any $m \in \ker \partial$ and any $P \in A_1$, the element $m \cdot P$ equals $m \cdot f(x)$ for some polynomial $f(x) \in k[x]$; and if this element is nonzero, then by applying a suitable power of ∂, one can recover m. Since $\ker \partial \cdot A_1$ is finitely generated over A_1, it follows that $\ker \partial$ must be finitely generated over k, hence finite-dimensional.

Bernstein's idea for the general case is to use an algebraic analogue of the Fourier transform. Recall that the usual Fourier transform (on functions) interchanges partial derivatives and multiplication by coordinate functions. We can imitate this algebraically by the following definition. Let M be a right A_n-module. Its *Fourier transform* is a left A_n-module $\hat{M}$, defined as follows: as a k-vector space, one has $\hat{M} = M$, but with A_n-action defined by

$$x_j \cdot m = m\partial_j \quad \text{and} \quad \partial_j \cdot m = mx_j.$$

To show that this gives $\hat{M}$ the structure of a left A_n-module, one has to check the relation $[\partial_i, x_j] = \delta_{i,j}$. This holds because

$$[\partial_i, x_j] \cdot m = \partial_i(x_j m) - x_j(\partial_i m) = m\partial_j x_i - mx_i\partial_j = m[\partial_j, x_i] = \delta_{i,j} m.$$

Its usefulness for studying direct images comes from the following lemma.

Lemma 19.4. *Consider a coordinate projection and its dual closed embedding*

$$\begin{aligned} p\colon \mathbb{A}^{n+1} \to \mathbb{A}^n, &\quad p(x_0, x_1, \dots, x_n) = (x_1, \dots, x_n), \\ i\colon \mathbb{A}^n \hookrightarrow \mathbb{A}^{n+1}, &\quad i(x_1, \dots, x_n) = (0, x_1, \dots, x_n). \end{aligned}$$

If M is a holonomic right A_{n+1}-module, then

$$H^j p_+ M \cong \widehat{L^j i^* \hat{M}}$$

for every $j \in \mathbb{Z}$.

Proof. By pretty much the same calculation that we did in Lecture 17, the direct image $p_+ M$ is computed by the relative version of the Spencer complex; in the case at hand, this is the complex of right A_n-modules

$$M \xrightarrow{\partial_0} M$$

Its cohomology lives in degree -1 and 0:

$$H^j p_+ M = \begin{cases} \ker(\partial_0 \colon M \to M) & \text{if } j = -1, \\ \operatorname{coker}(\partial_0 \colon M \to M) & \text{if } j = 0, \\ 0 & \text{otherwise.} \end{cases}$$

The right A_n-module structure on $H^j p_+ M$ is induced by the right A_{n+1}-module structure on M in the obvious way. On the other hand, the inverse image $\mathbf{L} i^* \hat{M}$ is computed by the complex of left A_n-modules

$$\hat{M} \xrightarrow{x_0} \hat{M}.$$

Its cohomology also lives in degree -1 and 0:

$$L^j i^* \hat{M} = \begin{cases} \ker(x_0 \colon \hat{M} \to \hat{M}) & \text{if } j = -1, \\ \operatorname{coker}(x_0 \colon \hat{M} \to \hat{M}) & \text{if } j = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Here the left A_n-module structure on $L^j i^* \hat{M}$ is induced by the left A_{n+1}-module structure on $\hat{M}$ in the obvious way. Since left multiplication by x_0 on $\hat{M}$ is, by definition, the same as right multiplication by ∂_0 on M, we have $H^j p_+ M = L^j i^* \hat{M}$ as k-vector spaces. The additional Fourier transform makes sure that the right A_n-module structures on both sides agree. □

The Fourier transform preserves holonomicity.

Lemma 19.5. *A right A_n-module M is holonomic if and only if its Fourier transform $\hat{M}$ is holonomic as a left A_n-module.*

Proof. We use the characterization of holonomicity in terms of Hilbert polynomials (from Lecture 3). Recall the definition of the Bernstein filtration

$$F_j^B A_n = \left\{ P = \sum c_{\alpha,\beta} x^\alpha \partial^\beta \;\middle|\; |\alpha| + |\beta| \le j \right\}.$$

If $F_\bullet M$ is a good filtration, compatible with the Bernstein filtration, then for $j \gg 0$, the function $j \mapsto \dim_k F_j M$ is a polynomial in j; the degree of this polynomial is denoted by $d(M)$. We showed in Lecture 6 that M is holonomic (in the sense that its characteristic variety has dimension n) if and only if $d(M) = n$. The proof of the lemma is now a triviality: simply observe that a good filtration $F_\bullet M$ is also a good filtration $F_\bullet \hat{M}$, due to the fact that the Bernstein filtration is symmetric in $x_1, \dots, x_n$ and $\partial_1, \dots, \partial_n$. It follows that $d(M) = d(\hat{M})$, and so M is holonomic iff $\hat{M}$ is holonomic. □

The last thing we need to check is that localization preserves holonomicity.

Lemma 19.6. *Let M be a holonomic left A_{n+1}-module. Then*

$$N = k[x_0, \dots, x_n, x_0^{-1}] \otimes_{k[x_0,\dots,x_n]} M$$

is again a holonomic A_{n+1}-module.

Proof. The argument is the same as in the proof of Proposition 3.10. We are going to make use of the numerical criterion for holonomicity in Lemma 3.11: Suppose that N is a left A_{n+1}-module, and $F_\bullet N$ a filtration compatible with the Bernstein filtration on A_{n+1}, such that

$$\dim_k F_j N \leq \frac{c}{(n+1)!} j^{n+1} + c_1 (j+1)^n$$

for some constants $c, c_1 \geq 1$. Then N is holonomic.

A suitable filtration on N is obtained by setting

$$F_j N = x_0^{-j} \otimes F_{2j} M$$

for every $j \geq 0$. It is easy to see that this filtration is compatible with the Bernstein filtration. Let us check that it is exhaustive. Any element of N can be written in the form $x_0^{-j} \otimes m$ for some $m \in M$ and some $j \geq 0$. Since $F_\bullet M$ is exhaustive, we have $m \in F_k M$ for some $k \geq 0$. Now

$$y_0^{-j} \otimes m = y_0^{-(j+\ell)} \otimes (y_0^\ell m),$$

and since $y_0^\ell m \in F_{k+\ell} M$, this element will belong to $F_{j+\ell} N$ as long as $k+\ell \leq 2(j+\ell)$ or, equivalently, as long as $\ell \geq k - 2j$.

Let us count dimensions. Since M is holonomic, we have $\dim_k F_j M = \chi(j)$, where $\chi(t) \in \mathbb{Q}[t]$ is a polynomial of degree $d(M) = n+1$. But then

$$\dim_k F_j N = \dim_k F_{2j} M = \chi(2j)$$

is still a polynomial of degree $n+1$; by the numerical criterion, this implies that N is again holonomic. □

Let us now put everything together and prove Theorem 18.5. By the argument from last time, it suffices to show that if M is a holonomic right A_{n+1}-module, and

$$p \colon \mathbb{A}^{n+1} \to \mathbb{A}^n, \quad p(x_0, x_1, \dots, x_n) = (x_1, \dots, x_n)$$

the coordinate projection, then $H^j p_+ M$ is holonomic for every $j \in \mathbb{Z}$. Let $\hat{M}$ be the Fourier transform of M; by Lemma 19.5, this is a holonomic left A_{n+1}-module. According to Lemma 19.4, we have

$$H^j p_+ M \cong \widehat{L^j i^* \hat{M}},$$

and so again by Lemma 19.5, it will be enough to prove that each $L^j i^* \hat{M}$ is a holonomic left A_n-module. By Lemma 19.1, the two potentially nonzero modules (for $j = -1$ and $j = 0$) are the kernel and cokernel of the morphism

$$\hat{M} \to k[x_0, \dots, x_n, x_0^{-1}] \otimes_{k[x_0,\dots,x_n]} \hat{M}.$$

The localization is again holonomic (by Lemma 19.6), and so the kernel and cokernel are holonomic modules. This suffices to conclude the proof.

Consequences. Let me point out a few interesting consequences of the result we have just proved.

First, consider the case where $f\colon X \to \operatorname{Spec} k$ is the morphism to a point. Given a holonomic right $\mathscr{D}_X$-module $\mathcal{M}$, the direct image $f_+\mathcal{M}$ is computed as the hypercohomology of the Spencer complex $\operatorname{Sp}(\mathcal{M})$. Thus Theorem 18.5 is saying that the hypercohomology of $\operatorname{Sp}(\mathcal{M})$ is a finite-dimensional k-vector space. In the special case $\mathcal{M} = \omega_X$, this says that the algebraic de Rham cohomology groups $H^j_{dR}(X/k)$ are finite-dimensional even if X is not proper. (When $k = \mathbb{C}$, this also follows from the isomorphism $H^j_{dR}(X/\mathbb{C}) \cong H^j(X,\mathbb{C})$ and some basic facts about the topology of nonsingular algebraic varietes.) One way to think about this is to consider the hypercohomology of $\operatorname{Sp}(\mathcal{M})$ as being something like the cohomology of X with coefficients in $\mathcal{M}$; the theorem is claiming that this cohomology is finite-dimensional whenever $\mathcal{M}$ is holonomic.

Second, consider the case of a closed embedding $i\colon Z \hookrightarrow X$. Here, the statement is that $\mathbf{L}i^*\mathcal{M}$ is holonomic for every holonomic left $\mathscr{D}_X$-module $\mathcal{M}$, even if $\mathcal{M}$ does not have the non-characteristic property. In particular, we can pull back along

$$i_x\colon \operatorname{Spec} k \hookrightarrow X$$

for any closed point $x \in X(k)$, and for any holonomic $\mathscr{D}_Y$-module $\mathcal{M}$, or any complex in $D^b_h(\mathscr{D}_Y)$, the inverse image $\mathbf{L}i_x^*\mathcal{M}$ is holonomic on $\operatorname{Spec} k$, hence has finite-dimensional cohomology. This is another important finiteness property of holonomic modules. It is obvious on the open subset where $\mathcal{M}$ is a vector bundle with integrable connection, but not at other points of Y.

Note. In fact, one can show that when k is algebraically closed, holonomic complexes are characterized by this finiteness property: an object $\mathcal{M}^\bullet \in D^b_{coh}(\mathscr{D}_X)$ belongs to the subcategory $D^b_h(\mathscr{D}_X)$ if, and only if, for every closed point $x \in X(k)$, the complex $\mathbf{L}i_x^*\mathcal{M}^\bullet$ has finite-dimensional cohomology. We don't have time to prove this, unfortunately.